

HB31
M415
no. 06-15
2009



DEWEY

DEWEY

Massachusetts Institute of Technology
Department of Economics
Working Paper Series

CAPITAL TAXATION

Quantitative Exploration of the Inverse Euler Equation

Emmanuel Farhi
Iván Werning

Working Paper 06-15
November 30, 2005
Revised version: April 10, 2009

Room E52-251
50 Memorial Drive
Cambridge, MA 02142

This paper can be downloaded without charge from the
Social Science Research Network Paper Collection at
<http://ssrn.com/abstract=902070>

Capital Taxation*

Quantitative Explorations of the Inverse Euler Equation

Emmanuel Farhi

Harvard University
efarhi@harvard.edu

Iván Werning

MIT
iwerning@mit.edu

April 10, 2009 (7:11am)

Abstract

This paper studies the efficiency gains from distorting savings in dynamic Mirrleesian private-information economies. We develop a method that perturbs the consumption process optimally, while preserving incentive compatibility. The Inverse Euler equation holds at the new optimized allocation. Starting from an equilibrium where agents can save freely allows us to compute the efficiency gains from savings distortions. We investigate how these gains depend on a limited set of features of the economy. We find an important role for general equilibrium effects. In particular, efficiency gains are greatly reduced when, rather than assuming a fixed interest rate, decreasing returns to capital are incorporated with a neoclassical technology. We compute the efficiency gains for the incomplete market model in Aiyagari [1994] and find them to be relatively modest for the baseline calibration. For higher levels of uncertainty, the efficiency gains can be sizable, but we find that most of the improvements can then be attributed to the relaxation of borrowing constraints, rather than the introduction of savings distortions.

* Farhi is grateful for the hospitality of the University of Chicago. Werning is grateful for the hospitality of Harvard University and the Federal Reserve Bank of Minneapolis during the time this draft was completed. We thank Florian Scheuer for outstanding research assistance. We thank comments and suggestions from Daron Acemoglu, Manuel Amador, Marios Angeletos, Patrick Kehoe, Narayana Kocherlakota, Mike Golosov, Ellen McGrattan, Robert Shimer, Aleh Tsyvinski and seminar participants at the Federal Reserve Bank of Minneapolis, Chicago, MIT, Harvard, Urbana-Champaign, Carnegie-Mellon, Wharton and the Conference on the Macroeconomics of Imperfect Risk Sharing at the UC-Santa Barbara, Society for Economic Dynamics in Vancouver and the NBER Summer institute. Special thanks to Mark Aguiar, Pierre-Olivier Gourinchas, Fatih Guvenen, Dirk Krueger, Fabrizio Perri, Luigi Pistaferri and Gianluca Violante for enlightening exchanges on the available empirical evidence for consumption and income. All remaining errors are our own.

Digitized by the Internet Archive
in 2011 with funding from
Boston Library Consortium Member Libraries

<http://www.archive.org/details/capitaltaxationq00farh>

1 Introduction

Recent work has upset a cornerstone result in optimal tax theory. According to Ramsey models, capital income should eventually go untaxed [Chamley, 1986, Judd, 1985]. In other words, individuals should be allowed to save freely and without distortions at the social rate of return to capital. This important benchmark has dominated formal thinking on this issue. By contrast, in economies with idiosyncratic uncertainty and private information it is generally suboptimal to allow individuals to save freely: constrained efficient allocations satisfy an *Inverse Euler equation*, instead of the agent's standard intertemporal Euler equation [Diamond and Mirrlees, 1977, Rogerson, 1985, Ligon, 1998]. Recently, extensions of this result have been interpreted as counterarguments to the Chamley-Judd no-distortion benchmark [Golosov, Kocherlakota, and Tsyvinski, 2003, Albanesi and Sleet, 2006, Kocherlakota, 2005, Werning, 2002].

This paper explores the quantitative importance of these arguments. To do so, we start from an equilibrium where individuals save freely, without distortions, and examine the efficiency gains obtained from introducing optimal savings distortions. The issue we address is largely unexplored because—deriving first-order conditions aside—it is difficult to solve dynamic economies with private information, except for some very particular cases, such as shocks that are i.i.d. over time.¹ We develop a new approach that sidesteps these difficulties.

We lay down an infinite-horizon Mirrleesian economy with neoclassical technology. Agents consume and work in every period. Preferences between consumption and work are assumed additively separable. Agents experience skill shocks that are private information, so that feasible allocations must be incentive-compatible.

Constrained efficient allocations satisfy a simple intertemporal condition for consumption known as the Inverse Euler equation. This condition is incompatible with the agents' standard Euler equation, implying that constrained efficient allocations cannot be decentralized in competitive equilibria where agents save freely at the technological rate of return to capital. Some form of savings distortion is needed.

Equivalently, starting from an incentive compatible allocation obtained from an equilibrium where agents save freely, efficiency gains are possible with the introduction of saving distortions. In this paper we are interested in computing these gains. We do so by perturbing the consumption assignment and holding the labor assignment unchanged, while preserving incentive compatibility. The new allocation satisfies the Inverse Euler equation and delivers the same utility while freeing up resources. The reduction in re-

¹ Other special cases that have been extensively explored are unemployment and disability insurance.

sources is our measure of efficiency gains. By leaving the labor assignment unchanged, we focus on the efficiency gains of introducing savings distortions, without changing the incentive structure, implicit in the labor assignment. In this way, we sidestep resolving the optimal trade-off between insurance and incentives.

There are several advantages to our approach. To begin with, our exercise does not require specifying some components of the economy. In particular, no knowledge of individual labor assignment or the disutility of work function is required. In this way, the degree to which work effort responds to incentives is not needed. This robustness is important, since empirical knowledge of these elasticities remains incomplete. Indeed, our efficiency gains depend only on the original consumption assignment, the utility function for consumption and technology. The planning problem we set up minimizes resources over a class of perturbations for consumption. This problem has the advantage of being tractable, even for rich specifications of uncertainty. In our view, having this flexibility is important for quantitative work. Furthermore, efficiency gains are shaped by intuitive properties involving both partial equilibrium and general equilibrium considerations, such as the variance of consumption growth and its dispersion across agents and time, the coefficient of relative risk aversion and the concavity of the production function. Finally, our measure of efficiency gains can provide lower and upper bounds on the full efficiency gains obtained from joint reforms that introduce savings distortions and change the assignment of labor.

Two approaches are possible to quantify the magnitude of these efficiency gains, and we pursue both. The first approach uses directly as a baseline allocation an empirically plausible process of individual consumption. The second approach recognizes that the empirical knowledge of consumption risk is more limited than that of income risk. It uses as a baseline allocation the outcome of a competitive economy where agents are subject to idiosyncratic skill shocks and can save freely by accumulating a risk-free bond.

Following the first approach, we analyze a model where the baseline allocation features geometric random walk processes for individual consumption and utility is logarithmic. In this case, we obtain closed form solutions. The perturbed allocation is simple. It is obtained by multiplying the baseline consumption assignment by a deterministic declining sequence. The size of this downward drift and the efficiency gains are increasing in the variance of consumption growth, which indexes the strength of the precautionary savings motive. Using a fixed interest rate, the efficiency gains span a wide range, going from 0% to 10%, depending primarily on the variance of consumption growth, for which empirical evidence is limited.

We find that, general equilibrium effects can dramatically mitigate these numbers.

This is because tilting individual consumption profiles requires decumulating capital. With a neoclassical technology, as capital is decumulated, the return to capital increases, raising the cost of further decumulation. With a capital share of $1/3$, we show that efficiency gains range from 0% to 0.25%, for plausible values of the variance in consumption growth.

Turning to the second approach, we adopt the incomplete-market specification from the seminal work of Aiyagari [1994]. In this economy, there is no aggregate uncertainty. Individuals face idiosyncratic labor income risk that they cannot insure. They can save in a risk-free asset, but cannot borrow. At a steady state equilibrium the interest rate is constant and equal to the marginal product of capital. Although individual consumption fluctuates, the cross-sectional distribution of assets and consumption is invariant. We take our baseline allocation from this steady state equilibrium.

Even with logarithmic utility, taking baseline consumption from this equilibrium model implies two differences relative to the geometric random-walk case, discussed previously. On the one hand, as is well known, agents are able to smooth consumption quite effectively in these Bewley models. This tends to minimize the variance of consumption growth and pushes towards low efficiency gains. On the other hand, consumption is not a geometric random walk. In particular, expected consumption growth is not equalized over time or across agents. Indeed, a steady state requires a stable cross-sectional distribution for consumption, implying some mean-reverting forces. This pushes for greater efficiency gains, because there are gains from aligning expected consumption growth in individual consumption across agents, without changing the growth in aggregate consumption or capital.²

We replicate the calibrations in Aiyagari [1994] and compute the steady state equilibrium for a number of income processes and values for the coefficient of relative risk aversion. For Aiyagari's baseline calibration, we find that efficiency gains are relatively small, below 0.2% for all utility specifications. Away from this baseline calibration, we find that efficiency gains increase with the coefficient of relative risk aversion and with the variance and persistence of the income process. Efficiency gains can be large if one combines high values of relative risk aversion with large and persistent shocks. However, for these calibrations we find that most of the efficiency gains are due to the relaxation of borrowing constraints, rather than to the introduction of savings distortions. These simulations illustrate our more general methodology and provide some insights into the

² As a result of equalizing expected consumption growth rates, the perturbed allocation will not feature an invariant distribution for consumption. Instead, the dispersion in the cross sectional distribution of consumption grows without bound. This is related to the "immiseration" result found in Atkeson and Lucas [1992].

determinants of the size of efficiency gains from savings distortions.

Related Literature. This paper relates to several strands of literature. First, there is the optimal taxation literature based on models with private information [see Golosov, Tsyvinski, and Werning, 2006, and the references therein]. Papers in this literature usually solve for the constrained efficient allocations, but few undertake a quantitative analysis of the efficiency gains due to savings distortions. Two exceptions are Golosov and Tsyvinski [2006] for disability insurance and Shimer and Werning [2008] for unemployment insurance. In both cases, the nature of the stochastic process for shocks allows for a low dimensional recursive formulation that is numerically tractable. Golosov and Tsyvinski [2006] provide a quantitative analysis of disability insurance. Disability is modeled as an absorbing negative skill shock. They calibrate their model and compute the welfare gains that can be reaped by moving from the most efficient allocation that satisfies free savings to the optimal allocation. They focus on logarithmic utility and report welfare gains of 0.5%. Shimer and Werning [2008] provide a quantitative analysis of unemployment insurance. They consider a sequential job search model, where a risk averse, infinitely lived worker samples wage offers from a known distribution. Regarding savings distortions, they show that with CARA utility allowing agents to save freely is optimal. With CRRA utility, savings distortions are optimal, but they find that the efficiency gains they provide are minuscule. As most quantitative exercises to date, both Golosov and Tsyvinski [2006] and Shimer and Werning [2008] are set in partial equilibrium settings with linear technologies.

Second, following the seminal paper by Aiyagari [1994], there is a vast literature on incomplete-market Bewley economies within the context of the neoclassical growth model. These papers emphasize the role of consumers self-smoothing through the precautionary accumulation of risk-free assets. In most positive analyses, government policy is either ignored or else a simple transfer and tax system is included and calibrated to current policies. In some normative analyses, some reforms of the transfer system, such as the income tax or social security, are evaluated numerically [e.g. Conesa and Krueger, 2005]. Our paper bridges the gap between the optimal-tax and incomplete-market literatures by evaluating the importance of the constrained-inefficiencies in the latter.

The notion of efficiency used in the present paper is often termed *constrained-efficiency*, because it imposes the incentive-compatibility constraints that arise from the assumed asymmetry of information. Within exogenously incomplete-market economies, a distinct notion of constrained-efficiency has emerged [see Geanakoplos and Polemarchakis, 1985]. The idea is roughly whether, taking the available asset structure as given, individuals could change their trading positions in such a way that generates a Pareto improvement

at the resulting market-clearing prices. This notion has been applied by Davila, Hong, Krusell, and Rios-Rull [2005] to Aiyagari's [1994] setup. They show that, in this sense, the resulting competitive equilibrium is inefficient. In this paper we also apply our methodology to examine an efficiency property of the equilibrium in Aiyagari's [1994] model, but it should be noted that our notion of constrained efficiency, which is based on preserving incentive-compatibility, is very different.

2 A Two-Period Economy with Linear Technology

We start with a simple two-period economy with linear technology and then extend the concepts to an infinite horizon setting with general technologies.

Preferences. There are two periods $t = 0, 1$. Agents are ex ante identical. We focus on symmetric allocations. Consumption takes place in both periods, while work occurs only in period $t = 1$. Agents obtain utility

$$v = U(c_0) + \beta \mathbb{E} [U(c_1) - V(n_1; \theta)]. \quad (1)$$

where U is the utility function from consumption, V is the disutility function from effective units of labor (hereafter: labor for short) and \mathbb{E} is the expectations operator.

Uncertainty is captured by an individual shock $\theta \in \Theta$ that affects the disutility of effective units of labor, where Θ is an interval of \mathbb{R} . We will sometimes refer to θ as a skill shock. To capture the idea that uncertainty is idiosyncratic, we assume that a version of the law of large number holds so that for any function f on Θ , $\mathbb{E}[f]$ corresponds to the average of f across agents.

The utility function U is assumed increasing, concave and continuously differentiable. We assume that the disutility function V is continuously differentiable and that, for any $\theta \in \Theta$, the function $V(\cdot, \theta)$ is increasing and convex. We also assume the single crossing property that $\frac{\partial}{\partial n_1} V(n_1; \theta)$ is strictly decreasing in θ , so that a high shock θ indicates a low disutility from work.

Incentive-Compatibility. The shock realizations are private information to the agent, so we must ensure that allocations are incentive compatible. By the revelation principle we can consider, without loss of generality, a direct mechanism where agents report their shock realization θ in period $t = 1$ and are assigned consumption and labor as a function of this report. The agent's strategy is a mapping $\sigma : \Theta \rightarrow \Theta$ where $\sigma(\theta)$ denotes the report made when the true shock is θ . The truth telling strategy is denoted by σ^* , which is defined by $\sigma^*(\theta) = \theta$ for all $\theta \in \Theta$.

It is convenient to change variables and define an allocation by the triplet $\{u_0, u_1, n_1\}$ with $u_0 \equiv U(c_0)$, $u_1(\theta) \equiv U(c_1(\theta))$. Incentive compatibility requires that truth-telling be optimal:

$$\sigma^* \in \arg \max_{\sigma} \{u_0 + \beta \mathbb{E} [u_1(\sigma(\theta)) - V(n_1(\sigma(\theta)); \theta)]\}.$$

Let θ^* be an arbitrary point in Θ . The single crossing property implies that incentive compatibility is equivalent to the condition that n_1 be non-decreasing and

$$u_1(\theta) - V(n_1(\theta), \theta) = u_1(\theta^*) - V(n_1(\theta^*), \theta^*) + \int_{\theta^*}^{\theta} \frac{\partial V}{\partial \hat{\theta}}(n_1(\hat{\theta}); \hat{\theta}) d\hat{\theta}. \quad (2)$$

Technology. We assume that technology is linear with labor productivity in period $t = 1$ equal to w and a rate of return on savings equal to q^{-1} . The resource constraints are

$$c(u_0) + k_1 \leq k_0 \quad (3)$$

$$\mathbb{E} [c(u_1(\theta)) - wn_1(\theta)] \leq q^{-1}k_1 \quad (4)$$

where $c \equiv U^{-1}$ is the inverse of the utility function.

For an allocation to satisfy both resource constraints with equality the required level of initial capital k_0 must be

$$k_0 = c(u_0) + q\mathbb{E} [c(u_1(\theta)) - wn_1(\theta)].$$

In what follows, we refer to k_0 as the cost of the allocation.

Perturbations. Given a utility level v and a non-decreasing labor assignment $\{n_1\}$, incentive compatibility and the requirement that the allocation deliver utility v determines a set of allocations $\Gamma(\{n_1\}, v)$. Indeed, this set has a simple structure. Substituting equation (2) into equation (1) implies that

$$u_0 + \beta u_1(\theta^*) = v + \beta V(n_1(\theta^*), \theta^*) - \beta \mathbb{E} \left[\int_{\theta^*}^{\theta} \frac{\partial V}{\partial \hat{\theta}}(n_1(\hat{\theta}); \hat{\theta}) d\hat{\theta} \right].$$

For given v and $\{n_1\}$, the right hand side is fixed. For any value of u_0 this equation can be seen as determining the value of $u_1(\theta^*)$. Equation (2) then determines the entire assignment $\{u_1\}$. Thus, elements of $\Gamma(\{n_1\}, v)$ are uniquely determined by the value of u_0 .

It is useful to restate this property as a perturbation. Given a baseline allocation

$\{u_0, u_1, n_1\} \in \Gamma(\{n_1\}, v)$, any other allocation $\{\tilde{u}_0, \tilde{u}_1, n_1\} \in \Gamma(\{n_1\}, v)$ satisfies

$$\tilde{u}_0 = u_0 - \beta\Delta \text{ and } \tilde{u}_1(\theta) = u_1(\theta) + \Delta \text{ for all } \theta \in \Theta,$$

for some $\Delta \in \mathbb{R}$. The reverse is also true, for any baseline allocation in $\Gamma(\{n_1\}, v)$ and any $\Delta \in \mathbb{R}$ a utility assignment constructed in this way is part of an allocation in $\Gamma(\{n_1\}, v)$.

Note that the construction of this perturbation is independent of the labor assignment $\{n_1\}$. To reflect this denote this corresponding set of utility assignments by $Y(\{u_0, u_1\}, 0) = \{\{u_0 - \beta\Delta, u_1 + \Delta\} \mid \Delta \in \mathbb{R}\}$. This notation captures the following point. For the purposes of describing all possible utility assignments consistent with a labor assignment $\{n_1\}$, knowledge of the labor assignment itself and the disutility function V can be replaced by knowledge of some baseline utility assignment.³

Free-Savings and Euler equation. The allocation $\{u_0, u_1, n_1\}$ is part of a free-savings equilibrium with natural borrowing limits if and only if truth telling and saving zero is optimal:

$$(\sigma^*, 0) \in \arg \max_{(\sigma, k)} \left\{ U(c(u_0) - k) + \beta \mathbb{E} \left[U(c(u_1(\sigma(\theta))) + q^{-1}k) - V(n_1(\sigma(\theta)); \theta) \right] \right\} \quad (5)$$

For short, we say that the allocation satisfies free-savings. Free-savings implies incentive compatibility and the Euler equation

$$U'(c(u_0)) = \beta q^{-1} \mathbb{E} [U'(c(u_1(\theta)))] \quad (6)$$

The converse is not generally true, these two conditions are not sufficient for free-savings. However, for a given labor assignment $\{n_1\}$ and utility level v , the utility assignments $\{u_0, u_1\}$ are uniquely determined by incentive compatibility, the requirement that the allocation deliver expected utility v and the Euler equation. That is, there exists a unique

³A note on our notation is in order. We model shocks as continuous, i.e. Θ is an interval of \mathbb{R} , but do not necessarily assume that π has full support on Θ . When it does not, our notation still requires defining the allocation for values of θ outside the support. Thus, we distinguish allocations that coincide on the support of π , but differ elsewhere. This is absolutely without loss of generality, but plays a role in the definitions of $\Gamma(\{n_1\}, v)$ and $Y(\{u_0, u_1\}, 0)$.

For example, consider the case where the support of π is composed of a finite set of points, so that θ belongs to a finite set with probability one. This ‘finite shock’ setting is often adopted as a simplifying assumption. In this case, one could define the allocation as a finite vector, with elements corresponding to the points in the support of π . However, with this alternative representation, there can exist two utility assignments, $\{u_0, u_1\}$ and $\{\tilde{u}_0, \tilde{u}_1\}$, that, together with $\{n_1\}$ are incentive compatible, but are not obtained from one another through the parallel perturbations we described. To see why this is consistent with the notation we adopt, note that in this case it is possible to define various labor assignments on Θ that coincide on the finite support of Θ . For each of these, our representation of all possible utility assignments $\Gamma(\{n_1\}, v)$ holds.

allocation in $\Gamma(\{n_1\}, v)$ that satisfies the Euler equation. For any given assignment $\{n_1\}$ the resulting allocation may or may not satisfy free savings. We denote by $D(v)$ the set of labor assignments that are compatible with free savings, given utility v .

Efficiency and The Inverse Euler Equation. Consider the allocation $\{u_0, u_1, n_1\}$ in $\Gamma(\{n_1\}, v)$ with the minimal cost - the most efficient allocation in $\Gamma(\{n_1\}, v)$. This allocation is uniquely determined by the requirement that $\{u_0, u_1, n_1\} \in \Gamma(\{n_1\}, v)$ and the following first order condition

$$\frac{1}{U'(c(u_0))} = \frac{1}{\beta q^{-1}} \mathbb{E} \left[\frac{1}{U'(c(u_1(\theta)))} \right]. \quad (7)$$

Equation (7) is the so called *Inverse Euler equation*.

It is useful to contrast equation (7) with the standard Euler equation (6). Jensen's inequality to equation (18) implies that at the optimum

$$U'(c(u_0)) < \beta q^{-1} \mathbb{E}[U'(c(u_1(\theta)))], \quad (8)$$

as long as consumption at $t + 1$ is uncertain condition on information at t . This shows that equation (7) is incompatible with equation (6). Thus, the minimal cost allocation $\{u_0, u_1, n_1\}$ in $\Gamma(\{n_1\}, v)$ cannot allow agents to save freely at the technology's rate of return, since then equation (6) would hold as a necessary condition, which is incompatible with the planner's optimality condition, equation (7).

For any allocation, define $\tau \in \mathbb{R}$ by

$$U'(c(u_0)) = q^{-1}(1 - \tau) \mathbb{E}[U'(c(u_1(\theta)))],$$

a measure of the distortion in the Euler equation that is sometimes referred to as the intertemporal wedge or implicit tax on savings.⁴ If equation (7) holds then $\tau > 0$.

Efficiency Gains. Given a utility level v and a labor assignment $\{n_1\}$, we define the following cost functions

$$\begin{aligned} \chi(\{n_1\}, v) &\equiv \min_{\{\tilde{u}_0, \tilde{u}_1, \tilde{n}_1\} \in \Gamma(\{n_1\}, v)} \{c(\tilde{u}_0) + q \mathbb{E}[c(\tilde{u}_1(\theta)) - w \tilde{n}_1(\theta)]\} \\ \chi^E(\{n_1\}, v) &\equiv \min_{\{\tilde{u}_0, \tilde{u}_1, \tilde{n}_1\} \in \Gamma(\{n_1\}, v)} \{c(\tilde{u}_0) + q \mathbb{E}[c(\tilde{u}_1(\theta)) - w \tilde{n}_1(\theta)]\} \quad \text{s.t. equation (6)} \end{aligned}$$

The first minimization implies that the corresponding allocation satisfies the Inverse Eu-

⁴We do not concern ourselves here with explicit tax systems that implement efficient allocations.




Figure 1: On the bottom left panel, the labor assignments corresponding to the minimum of the upper frontier is in $D(v)$. On the top right panel, the labor assignment corresponding to the minimum of the lower frontier is in $D(v)$. On the top left panel, both labor assignments are in $D(v)$. On the bottom right panel, none of these labor assignments is in $D(v)$.

ler equation. Given the discussion above, the second minimization takes place over a singleton—the allocation is determined by the constraints. By definition, it satisfies the Euler equation. Thus, the allocations behind these two costs satisfy, respectively, the Inverse Euler and standard Euler equations.

The Euler equation acts as an additional constraint in the second minimization. Therefore, the difference between these cost functions

$$\chi^E(\{n_1\}, v) - \chi(\{n_1\}, v) \quad (9)$$

is positive and captures the efficiency cost of imposing the Euler equation. Whenever $\{n_1\} \in D(v)$, this coincides with the cost of imposing free savings. This cost difference

(9) can hence also be interpreted as a measure of the efficiency gains from optimal savings distortions, and is the main focus of the paper. It has a number of desirable properties.

First, our measure of efficiency gains can be computed by a simple perturbation method. Indeed, consider the allocation $\{u_0, u_1, n_1\} \in \Gamma(\{n_1\}, v)$ that satisfies the Euler equation. By definition, its cost is $\chi^E(\{n_1\}, v)$. As discussed above, allocations in $\Gamma(\{n_1\}, v)$ can be obtained through simple perturbations $Y(\{u_0, u_1\}, 0) = \{\{u_0 - \beta\Delta, u_1 + \Delta\} \mid \Delta \in \mathbb{R}\}$ of the baseline utility assignment $\{u_0, u_1\}$ while fixing the labor assignment $\{n_1\}$. Within this class of perturbations, there exists a unique perturbed utility assignment $\{\tilde{u}_0, \tilde{u}_1\}$ that satisfies the Inverse Euler equation. Efficiency gains (9) are given by

$$\chi^E(\{n_1\}, v) - \chi(\{n_1\}, v) = c(u_0) - c(\tilde{u}_0) + q\mathbb{E}[c(u_1(\theta)) - c(\tilde{u}_1)]. \quad (10)$$

Second, given the baseline utility assignment $\{u_0, u_1\}$, no knowledge of either the labor assignment $\{n_1\}$ or the disutility of work V is needed in order to compute our measure of efficiency gains. In other words, one does not need to take a stand on how elastic work effort is to changes in incentives. More generally one does not need to take a stand on whether the problem is one of private information regarding skills or of moral hazard regarding effort, etc. This robustness is a crucial advantage since current empirical knowledge of these characteristics and parameters is limited and controversial. The efficiency gains (9) address precisely the question of whether the intertemporal allocation of consumption is efficient, without taking a stand on how correctly tradeoff between insurance and incentives has been resolved.⁵

Third, as this paper will show, the magnitude of the efficiency gains (9) is determined by some rather intuitive properties involving both partial equilibrium and general equilibrium considerations: the relative variance of consumption changes, the coefficient of relative risk aversion and the concavity of the production function.

Finally, we now show how (9) is informative about more encompassing measures of efficiency gains that also allow for changes in the labor assignment $\{n_1\}$.

Decomposition and Bounds. Suppose the economy is initially constrained by free savings. Suppose further that, subject to this restriction, the allocation is efficient. This is

⁵The robustness properties of the allocations in $\Gamma(\{n_1\}, v)$, and hence of our measure of efficiency gains, can be more formally described as follows. Consider the set $\Omega(\{u_0, u_1, n_1\}, v)$ of disutility functions V with the following properties: V is continuously differentiable; for any $\theta \in \Theta$, the function $V(\cdot, \theta)$ is increasing and convex; V has the single crossing property that $\frac{\partial}{\partial n_1} V(n_1; \theta)$ is strictly decreasing in θ and equation (2) holds. The set of allocations $\Gamma(\{n_1\}, v)$ is the largest of allocations such that for all $V \in \Omega(\{u_0, u_1, n_1\}, v)$, incentive compatibility (2) and promise keeping (1) hold.

represented by point A in the figure, with labor assignment

$$\{n_1^A\} \in \arg \min_{\{n_1\}} \chi^E(\{n_1\}, v) \quad \text{s.t. } \{n_1\} \in D(v).$$

Now consider lifting the restriction of free savings and allowing the optimal savings distortions, as described by the Inverse Euler equation. The new efficient allocation is represented by point C, with labor assignment $\{n_1^C\} \in \arg \min_{\{n_1\}} \chi(\{n_1\}, v)$. The move from A to C lowers costs by

$$\chi^E(\{n_1^A\}, v) - \chi(\{n_1^C\}, v) = \chi^E(\{n_1^A\}, v) - \chi^E(\{n_1^B\}, v) + \chi^E(\{n_1^B\}, v) - \chi(\{n_1^C\}, v).$$

The first term $\chi^E(\{n_1^A\}, v) - \chi^E(\{n_1^B\}, v)$ captures the move from point A to B, where $\{n_1^B\} \in \arg \min_{\{n_1\}} \chi^E(\{n_1\}, v)$. It represents the potential gains from removing the constraint $\{n_1\} \in D(v)$, but maintaining the Euler equation. The second term, $\chi^E(\{n_1^B\}, v) - \chi(\{n_1^C\}, v)$, captures the move from point B to point C. It represents the gains obtained from allowing optimal savings distortions, so that the Inverse Euler equation may be satisfied.

This decomposition emphasizes that savings distortions may be valuable for two qualitatively distinct reasons. First, the set of implementable labor assignments is enlarged. Second, for any labor assignment, perturbing the consumption assignment to satisfy the Inverse Euler equation, as opposed to the Euler equation, reduces costs. To the best of our knowledge, this decomposition is novel.

Note that the first source of efficiency gains may not be present. This is the case whenever $\{n_1^B\} \in D(v)$, so that points A and B coincide. These cases are shown in the top left and bottom left panels of the figure. Indeed, in numerical simulations we found that this is the typical case for standard utility functions and distributions π . In particular, we adopted iso-elastic utility and disutility functions $U(c) = c^{1-\sigma}/(1-\sigma)$ and $V(n; \theta) = \alpha(n/\theta)^\gamma$. As for π , we tried various classes of continuous distribution, including uniform, log-normal, exponential, and Pareto. For a given specification and parameters, we first computed the analogue of point B, by solving a planning problem that minimized cost subject to the Euler equation, incentive compatibility and promise keeping.⁶ We then verified that this allocation was compatible with free savings, so that $\{n_1^B\} \in D(v)$, by verifying condition (5) directly. We tried a wide range of preference and distribution parameters for these classes. In all cases, we found that $\{n_1^B\} \in D(v)$, so that the first source of efficiency gains was found to be nil.

⁶This can also be thought of as a “first order approach” for solving point A. In the next step, we check that A and B coincide, verifying the validity of the first order approach.

Of course, there are examples where this source of gains is positive. One such example is the case with finite shocks. However, our numerical findings with standard continuous distributions, where the gains are nil, suggest that conclusions based on finite shock examples, should be interpreted with caution. More work is needed to understand precisely the situations where this source of gains is nonzero, to determine the plausibility of these scenarios. In any case, this first source of efficiency gains is not the focus of this paper and we will not explore this issue further.

In contrast to the first source, the second source of efficiency gains is always present, since $\chi^E(\{n_1\}, v) - \chi(\{n_1\}, v) > 0$ as long as consumption is uncertain, so that the Inverse and Euler equation and the Euler equation are incompatible.

This two-period, linear technology example is tractable. In particular, points A, B and C can be computed numerically. However, in the rest of the paper we are interested in environments with concave technologies and an infinite horizon, with productivity modeled as a stochastic process. For these dynamic extensions, computing the analogues of points A, B and C is not tractable, except for some special cases such as i.i.d. or fully persistent shocks. However, note that

$$\chi^E(\{n_1^B\}, v) - \chi(\{n_1^B\}, v) \leq \chi^E(\{n_1^B\}, v) - \chi(\{n_1^C\}, v) \leq \chi^E(\{n_1^C\}, v) - \chi(\{n_1^C\}, v).$$

This shows that knowledge of the vertical distance between the two cost functions $\chi^E(\{n_1\}, v)$ and $\chi(\{n_1\}, v)$ is informative about the second source of efficiency gains.

3 Infinite Horizon

In this section, we lay down our general environment. We then describe a class perturbations that preserve incentive compatibility. These perturbations serve as the basis for our method to compute efficiency gains.

3.1 The Environment

We cast our model within a general Mirrleesian dynamic economy. Our formulation is closest to Golosov, Kocherlakota, and Tsyvinski [2003]. This paper obtains the Inverse Euler equation [e.g. Diamond and Mirrlees, 1977] in a general dynamic Mirrlees economy, where agents' privately observed skills evolve as a stochastic process.

Preferences. Our economy is populated by a continuum of agent types indexed by $i \in I$ distributed according to the measure ψ . Preferences generalize those used in Section 2

and are summarized by the expected discounted utility

$$\sum_{t=0}^{\infty} \beta^t \mathbb{E}^i [U(c_t^i) - V(n_t^i; \theta_t^i)]$$

\mathbb{E}^i is the expectations operator for type i .

Additive separability between consumption and leisure is a feature of preferences that we adopt because it is required for the arguments leading to the Inverse Euler equation.⁷

Idiosyncratic uncertainty is captured by an individual specific shock $\theta_t^i \in \Theta$, where as in Section 2, Θ is an interval of the real line. These shocks affect the disutility of effective units of labor. We sometimes refer to them as skill shocks. The stochastic process for each individual θ_t^i is identically distributed within each type $i \in I$ and independently distributed across all agents. We denote the history up to period t by $\theta^{i,t} \equiv (\theta_0^i, \theta_1^i, \dots, \theta_t^i)$, and by π^i the probability measure on Θ^∞ corresponding to the law of the stochastic process θ_t^i for an agent of type i .

Given any function f on Θ^∞ , we denote the integral $\int f(\theta^{i,\infty}) d\pi(\theta^{i,\infty})$ using the expectation notation $\mathbb{E}^i[f(\theta^{i,\infty})]$ or simply $\mathbb{E}^i[f]$. Similarly, we write $\mathbb{E}^i[f(\theta^{i,\infty})|\theta^{i,t-1}]$, or simply $\mathbb{E}_{t-1}^i[f]$, for the conditional expectation of f given history $\theta^{i,t-1} \in \Theta^t$.

As in Section 2, all uncertainty is idiosyncratic and we assume that a version of the law of large number holds so that for any function f on Θ^∞ , $\mathbb{E}^i[f]$ corresponds to the average of f across agents with type i .

To preview the use we will have for types $i \in I$, note that in our numerical implementation we will assume that skills follow a Markov process. We will consider allocations that result from a market equilibrium where agents save in a riskless asset. For this kind of economy, agent types are then initial asset holdings together with initial skill. The measure ψ captures the joint distribution of these two variables.

It is convenient to change variables, translating consumption allocation into utility assignments $\{u_t^i(\theta^{i,t})\}$, where $u_t^i(\theta^{i,t}) \equiv U(c_t^i(\theta^{i,t}))$. This change of variable will make incentive constraints linear and render the planning problem that we will introduce shortly convex. Then, agents of type $i \in I$ with allocation $\{u_t^i\}$ and $\{n_t^i\}$ obtain utility

$$v^i \equiv \sum_{t=0}^{\infty} \beta^t \mathbb{E}^i \left[u_t^i(\theta^{i,t}) - V(n_t^i(\theta^{i,t}); \theta_t^i) \right] \quad (11)$$

Information and Incentives. The shock realizations are private information to the agent.

⁷ The intertemporal additive separability of consumption also plays a role. However, the intertemporal additive separability of work effort is completely immaterial: we could replace $\sum_{t=0}^{\infty} \beta^t \mathbb{E}[V(n_t; \theta_t)]$ with some general disutility function $\hat{V}(\{n_t\})$.

We invoke the revelation principle to derive the incentive constraints by considering a direct mechanism. Agents are allocated consumption and labor as a function of the entire history of reports. The agent's strategy determines a report for each period as a function of the history, $\{\sigma_t^i(\theta^{i,t})\}$. The incentive compatibility constraint requires that truth-telling, $\sigma_t^{i,*}(\theta^{i,t}) = \theta_t^i$, be optimal:

$$\sum_{t=0}^{\infty} \beta^t \mathbb{E}^i [u_t^i(\theta^{i,t}) - V(n_t^i(\theta^{i,t}); \theta_t^i)] \geq \sum_{t=0}^{\infty} \beta^t \mathbb{E}^i [u_t^i(\sigma_t^{i,t}(\theta^{i,t})) - V(n_t^i(\sigma_t^{i,t}(\theta^{i,t})); \theta_t^i)] \quad (12)$$

for all reporting strategies $\{\sigma_t^i\}$ and all $i \in I$.

Technology. Let C_t and N_t represent aggregate capital, labor and consumption for period t , respectively. That is, letting $c \equiv U^{-1}$ denote the inverse of the utility function,

$$C_t \equiv \int \mathbb{E}^i [c(u_t^i)] d\psi$$

$$N_t \equiv \int \mathbb{E}^i [n_t^i] d\psi$$

for $t = 0, 1, \dots$. In order to facilitate our efficiency gains calculations, it will prove convenient to index the resource constraints by e_t which represents the aggregate amount of resources that is being economized in every period. The resource constraints are then

$$K_{t+1} + C_t + e_t \leq (1 - \delta)K_t + F(K_t, N_t) \quad t = 0, 1, \dots \quad (13)$$

where K_t denotes aggregate capital. The function $F(K, N)$ is assumed to be homogenous of degree one, concave and continuously differentiable, increasing in K and N .

Two cases are of particular interest. The first is the neoclassical growth model, where $F(K, N)$ is strictly concave and satisfies Inada conditions $F_K(0, N) = \infty$ and $F_K(\infty, N) = 0$. In this case, we also impose $K_t \geq 0$. The second case has linear technology $F(K, N) = N + (q^{-1} - 1)K$, $\delta = 0$ and $0 \leq q < 1$. One interpretation is that output is linear in labor with productivity normalized to one, and a linear storage technology with safe gross rate of return q^{-1} is available. Another interpretation is that this represents the economy-wide budget set for a partial equilibrium analysis, with constant interest rate $1 + r = q^{-1}$ and unit wage. Under either interpretation, we avoid corners by allowing capital to be negative up to the present value of future output, $K_{t+1} \geq -\sum_{s=1}^{\infty} q^s N_{t+s}$. This represents the natural borrowing limit and allows us to summarize the constraints on the economy

by the single present-value condition

$$\sum_{t=0}^{\infty} q^t C_t \leq \sum_{t=0}^{\infty} q^t (N_t - e_t) + \frac{1}{q} K_0.$$

Moral Hazard Extension. We could extend our framework to incorporate moral hazard in addition to private information. For example, each period agents could choose an unobservable action e_t^i that creates additively separable disutility. Effective labor, n_t^i , is a function of the history of effort and shocks $n_t^i = f_t(e^{i,t}, \theta^{i,t})$. In addition, the distribution over skill shocks may be affected by the sequence of effort, so that the distribution of $\theta^{i,t}$ depends on the effort choices $e^{i,t-1}$.

Although we will not pursue in further detail the notation required to formalize this type of model, all our subsequent analysis extends to such a setting. Indeed, our numerical applications may also be interpreted this way. This is important to keep in mind, since a hybrid model like this one seems more realistic than a model focusing exclusively on private information, as in the standard Mirrlees model.

Feasibility. An allocation $\{u_t^i, n_t^i, K_t, e_t\}$ and utility profile $\{v^i\}$ is *feasible* if conditions (11)–(13) hold. That is, feasible allocations that deliver utility v^i to agent of type $i \in I$, must be incentive compatible and resource feasible.

Free Savings. For the purposes of this paper, an important benchmark is the case where agents can save, and perhaps also borrow, freely. This increases the choices available to agents, which adds further restrictions relative to the incentive compatibility constraints.

In this scenario, the government enforces labor and taxes as a function of the history of reports, but does not control consumption directly. Disposable after-tax income is $w_t n_t^i(\sigma^{i,t}(\theta^{i,t})) - T_t^i(\sigma^{i,t}(\theta^{i,t}))$.⁸ Agents face the following sequence of budget and borrowing constraints:

$$c_t^i(\theta^{i,t}) + a_{t+1}^i(\theta^{i,t}) \leq w_t n_t^i(\sigma^{i,t}(\theta^{i,t})) - T_t^i(\sigma^{i,t}(\theta^{i,t})) + (1 + r_t) a_t^i(\theta^{i,t-1}) \quad (14a)$$

$$a_{t+1}^i(\theta^{i,t}) \geq \underline{a}_{t+1}^i(\sigma^{i,t}(\theta^{i,t})) \quad (14b)$$

with a_0^i given. We allow the borrowing limits $\underline{a}_{t+1}^i(\theta^{i,t})$ to be tighter than the natural borrowing limits.

Agents with type i maximize utility $\sum_{t=0}^{\infty} \beta^t \mathbb{E}^i[u_t^i(\sigma^{i,t}(\theta^{i,t})) - V(n_t^i(\sigma^{i,t}(\theta^{i,t})); \theta_t^i)]$ by

⁸A special case of interest is where the dependence of the tax on any history of reports $\hat{\theta}^{i,t}$, can be expressed through its effect on the history of labor $n^{i,t}(\hat{\theta}^{i,t})$. That is, when $T_t^i(\hat{\theta}^{i,t}) = T_t^{i,n}(n^{i,t}(\hat{\theta}^{i,t}))$ for some $T_t^{i,n}$ function.

choosing a reporting, consumption and saving strategy $\{\sigma_t^i, c_t^i, a_{t+1}^i\}$ subject to the sequence of constraints (14), taking a_0^i and $\{n_t^i, T_t^i, w_t, r_t\}$ as given. A feasible allocation $\{u_t^i, n_t^i, K_t, e_t\}$ is part of a free-savings equilibrium if there exist taxes $\{T_t^i\}$, such that the optimum $\{\sigma_t^i, c_t^i, a_{t+1}^i\}$ for agent of type i with wages and interest rates given by $w_t = F_N(K_t, N_t)$ and $r_t = F_K(K_t, N_t) - \delta$ satisfies truth telling $\sigma_t^i(\theta^{i,t}) = \theta_t^i$ and generates the utility assignment $u_t^i(\theta^{i,t}) = U(c_t^i(\theta^{i,t}))$.⁹

At a free-savings equilibrium, the incentive compatibility constraints (12) are satisfied. The consumption-savings choices of agents impose additional further restrictions. In particular, a necessary condition is the intertemporal Euler condition

$$U'(c(u_t^i)) \geq \beta(1 + r_{t+1})\mathbb{E}_t^i[U'(c(u_{t+1}^i))], \quad (15)$$

with equality if $a_{t+1}^i(\theta^{i,t}) > \underline{a}_{t+1}^i(\theta^{i,t})$. Note that if the borrowing limits $\underline{a}_{t+1}^i(\theta^{i,t})$ are equal to the natural borrowing limits, then the Euler equation (15) always holds with an equality.

Efficiency. We say that the allocation $\{u_t^i, n_t^i, K_t, e_t\}$ and utility profile $\{v^i\}$ is dominated by the alternative $\{\tilde{u}_t^i, \tilde{n}_t^i, \tilde{K}_t, \tilde{e}_t\}$ and $\{\tilde{v}^i\}$, if $\tilde{v}^i \geq v^i$, $\tilde{K}_0 \leq K_0$, $e_t \leq \tilde{e}_t$ for all periods t and either $\tilde{v}^i > v^i$ for a set of agent types of positive measure, $\tilde{K}_0 < K_0$ or $e_t < \tilde{e}_t$ for some period t .

We say that a feasible allocation is *efficient* if it is not dominated by any feasible allocation. We say that an allocation is *conditionally efficient* if it is not dominated by a feasible allocation with the same labor allocation $n_t^i = \tilde{n}_t^i$.

As explained in Section 2, allocations that are part of a free-savings equilibrium are not conditionally efficient. Conditionally efficient allocations satisfy a first order condition, the Inverse Euler equation, which is inconsistent with the Euler equation. Being part of a free-savings equilibrium therefore acts as a constraint on the optimal provision of incentives and insurance. Efficiency gains can be reaped by departing from free-savings.

3.2 Incentive Compatible Perturbations

In this section, we develop a class of perturbations of the allocation of consumption that preserve incentive compatibility. We then introduce a concept of efficiency, Δ -efficiency, that corresponds to the optimal use of these perturbations. Our perturbation set is large enough to ensure that every Δ -efficient allocation satisfies the Inverse Euler equation.

⁹Note that individual asset holdings are not part of this definition. This is convenient because asset holdings and taxes are indeterminate due to the usual Ricardian equivalence argument.

Moreover, we show that Δ -efficiency and conditional efficiency are closely related concepts: Δ -efficiency coincides with conditional efficiency on allocations that satisfy some mild regularity conditions. Finally, we derive some properties of Δ -efficient allocations with constant aggregates, which we term steady-states, in the case where the utility function is of the CRRA form.

A Class of Perturbations. For any period t and history $\theta^{i,t}$ a feasible perturbation, of any baseline allocation, is to decrease utility at this node by $\beta\Delta^i$ and compensate by increasing utility by Δ^i in the next period for all realizations of θ_{t+1}^i . Total lifetime utility is unchanged. Moreover, since only parallel shifts in utility are involved, incentive compatibility of the new allocation is preserved. We can represent the new allocation as $\tilde{u}_t^i(\theta^{i,t}) = u_t^i(\theta^{i,t}) - \beta\Delta^i$, $\tilde{u}_{t+1}^i(\theta^{i,t+1}) = u_{t+1}^i(\theta^{i,t+1}) + \Delta^i$, for all θ_{t+1}^i .

This perturbation changes the allocation in periods t and $t+1$ after history $\theta^{i,t}$ only. The full set of variations generalizes this idea by allowing perturbations of this kind at all nodes:

$$\tilde{u}_t^i(\theta^t) \equiv u_t^i(\theta^t) + \Delta^i(\theta^{i,t-1}) - \beta\Delta^i(\theta^{i,t})$$

for all sequences of $\{\Delta^i(\theta^{i,t})\}$ such that $\tilde{u}_t^i(\theta^{i,t}) \in U(\mathbb{R}_+)$ and such that the limiting condition

$$\lim_{T \rightarrow \infty} \beta^T \mathbb{E}^i[\Delta^i(\sigma^{i,T}(\theta^{i,T}))] = 0$$

for all reporting strategies $\{\sigma_t^i\}$. This condition rules out Ponzi-like schemes in utility.¹⁰ By construction, the agent's expected utility, for *any* strategy $\{\sigma_t^i\}$, is only changed by a constant Δ_{-1}^i :

$$\sum_{t=0}^{\infty} \beta^t \mathbb{E}^i[\tilde{u}_t^i(\sigma^{i,t}(\theta^{i,t}))] = \sum_{t=0}^{\infty} \beta^t \mathbb{E}^i[u_t^i(\sigma^{i,t}(\theta^{i,t}))] + \Delta_{-1}^i. \quad (16)$$

It follows directly from equation (12) that the baseline allocation $\{u_t^i\}$ is incentive compatible if and only if the new allocation $\{\tilde{u}_t^i\}$ is incentive compatible. Note that the value of the initial shifter Δ_{-1}^i determines the lifetime utility of the new allocation relative to its baseline. Indeed, for any fixed infinite history $\bar{\theta}^{i,\infty}$ equation (16) implies that (by substituting the deterministic strategy $\sigma_t^i(\theta^{i,t}) = \bar{\theta}_t^i$)

$$\sum_{t=0}^{\infty} \beta^t \tilde{u}_t^i(\bar{\theta}^{i,t}) = \sum_{t=0}^{\infty} \beta^t u_t^i(\bar{\theta}^{i,t}) + \Delta_{-1}^i \quad \forall \bar{\theta}^{i,\infty} \in \Theta^\infty \quad (17)$$

¹⁰ Note that the limiting condition is trivially satisfied for all variations with finite horizon: sequences for $\{\Delta_t^i\}$ that are zero after some period T . This was the case in the discussion of a perturbation at a single node and its successors.

Thus, ex-post realized utility is the same along all possible realizations for the shocks.¹¹

Let $Y(\{u_t^i\}, \Delta_{-1}^i)$ denote the set of utility allocations $\{\tilde{u}_t^i\}$ that can be generated by these perturbations starting from a baseline allocation $\{u_t^i\}$ for a given initial Δ_{-1}^i . This is a convex set.

Below, we show that these perturbation are rich enough to deliver the Inverse Euler equation. In this sense, they fully capture the characterization of optimality stressed by Golosov et al. [2003].

An allocation $\{u_t^i, n_t^i, K_t, e_t\}$ with utility profile $\{v^i\}$ is Δ -efficient if it is feasible and not dominated by another feasible allocation $\{\tilde{u}_t^i, \tilde{n}_t^i, \tilde{K}_t, \tilde{e}_t\}$ such that $\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, \Delta_{-1}^i)$.

Note that conditional efficiency implies Δ -efficiency, since both concepts do not allow for changes in the labor allocation. Under mild regularity conditions, the converse is also true. More precisely, in Appendix A, we define the notion of *regular* utility and labor assignments $\{u_t^i, n_t^i\}$.¹² We then show that Δ -efficiency coincides with conditional efficiency on the class of allocations with regular utility and labor assignments. Indeed, given a regular utility and labor assignment $\{u_t^i, n_t^i\}$, the perturbations $Y(\{u_t^i\}, \Delta_{-1}^i)$ characterize all the utility assignments $\{\tilde{u}_t^i\}$ such that $\{\tilde{u}_t^i, n_t^i\}$ is regular and satisfies the incentive compatibility constraints (12).

We will shortly present a methodology to compute the efficiency gains from restoring Δ -efficiency. We refer the reader to the discussion in Section 2 for an extensive motivation of this question.

Inverse Euler equation. Building on Section 2, we review briefly the Inverse Euler equation which is the optimality condition for any Δ -efficient allocation.

Proposition 1. *A set of necessary and sufficient conditions for an allocation $\{u_t^i, n_t^i, K_t, e_t\}$ to be Δ -efficient is given by*

$$c'(u_t^i) = \frac{q_t}{\beta} \mathbb{E}_t^i[c'(u_{t+1}^i)] \iff \frac{1}{U'(c(u_t^i))} = \frac{q_t}{\beta} \mathbb{E}_t^i \left[\frac{1}{U'(c(u_{t+1}^i))} \right]. \quad (18)$$

where $q_t = 1/(1 + r_t)$ and

$$r_t \equiv F_K(K_{t+1}, N_t) - \delta \quad (19)$$

is the technological rate of return.

¹¹ The converse is nearly true: by taking appropriate expectations of equation (17) one can deduce equation (16), but for a technical caveat involving the possibility of inverting the order of the expectations operator and the infinite sum (which is always possible in a version with finite horizon and Θ finite). This caveat is the only difference between equation (16) and equation (17).

¹² Regularity is a mild technical assumption which is necessary to derive an Envelope condition crucial for our proof.

A Δ -efficient allocation cannot allow agents to save freely at the technology's rate of return, since then equation (15) would hold as a necessary condition, which is incompatible with the planner's optimality condition, equation (18).

Another implication of the Inverse Euler equation (18) concerns the interest rates that prevail at steady states for Δ -efficient allocations. We refer to allocations $\{u_t^i, n_t^i, K_t, e_t\}$ with constant aggregates $C_t = C_{ss}$, $N_t = N_{ss}$, $K_t = K_{ss}$ and $e_t = e_{ss}$ as steady states. At a steady state, the discount factor is constant $q_t = q_{ss} \equiv 1/(1 + r_{ss})$ where $r_{ss} = F_K(K_{ss}, N_{ss}) - \delta$. Note that our notion of a steady state is only a condition on aggregate variables. It does not presume that individual consumption is constant nor that there is an invariant distribution for the cross section of consumption.

Suppose the utility function is of the constant relative risk aversion (CRRA) form, so that $U(c) = c^{1-\sigma}/(1-\sigma)$ for some $\sigma > 0$. The Inverse Euler equation is then $\mathbb{E}_t^i[c(u_{t+1}^i)^\sigma] = (\beta/q_{ss})c(u_t^i)^\sigma$. When $\sigma = 1$, this implies $C_{t+1} = (\beta/q_{ss})C_t$, so that aggregate consumption is constant if and only if $\beta = q_{ss}$. For $\sigma > 1$, Jensen's inequality implies that $C_{t+1} \leq (\beta/q_{ss})^{1/\sigma}C_t$, so that $\beta < q_{ss}$ is inconsistent with a steady state. The argument for $\sigma < 1$ is symmetric.

Proposition 2. *Suppose CRRA preferences $U(c) = c^{1-\sigma}/(1-\sigma)$ with $\sigma > 0$. Then at a steady state for a Δ -efficient allocation*

- (i) *if $\sigma \geq 1$ then $q_{ss} \leq \beta$ and $r_{ss} \geq \beta^{-1} - 1$*
- (ii) *if $\sigma \leq 1$ then $q_{ss} \geq \beta$ and $r_{ss} \leq \beta^{-1} - 1$*

This result can be contrasted with the properties of steady state equilibria in incomplete markets models, where agents are allowed to save freely at the technological rate of return. As shown by Aiyagari [1994], in such cases, the steady state interest rate is always lower than the discount rate $1/\beta - 1$.

4 Efficiency Gains

In this section we consider a baseline allocation and consider an improvement on it that yields a Δ -efficient allocation. We define a metric for efficiency gains from this improvement. This leads us to a planning problem to compute these gains. We then show how to solve this planning problem. In particular, we work with a relaxed planning problem that allows us to break up the solution into component planning problems. Moreover, we show how these problems may admit a recursive representation. Indeed, in the logarithmic case we find a closed-form solution. More generally, we show that each component

planning problem is mathematically isomorphic to solving an income fluctuations problem where Δ_t plays the role of wealth.

4.1 Planning Problem

If an allocation $\{u_t^i, n_t^i, K_t, e_t\}$ with corresponding utility profile $\{v^i\}$, is not Δ -efficient, then we can always find an alternative allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{e}_t\}$ that leaves utility unchanged, so that $\tilde{v}^i = v^i$ for all $i \in I$, but economizes on resources: $\tilde{K}_0 \leq K_0$ and $e_t \leq \tilde{e}_t$ with at least one strict inequality. In the rest of the paper, we restrict to cases where $e_t = 0$, $\tilde{K}_0 = K_0$ and $\tilde{e}_t = \tilde{\lambda} C_t$ for some $\tilde{\lambda} > 0$. We then take $\tilde{\lambda}$ as our measure of efficiency gains between these allocations. This measure represents the resources that can be saved in all periods in proportion to aggregate consumption.

We now introduce a planning problem that uses this metric to compute the distance of any baseline allocation from the Δ -efficient frontier. For any given baseline allocation $\{u_t^i, n_t^i, K_t, 0\}$, which is feasible with $e_t = 0$ for all $t \geq 0$, we seek to maximize $\tilde{\lambda}$ by finding an alternative allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ with $\tilde{K}_0 = K_0$,

$$\tilde{K}_{t+1} + \int \mathbb{E}^i[c(\tilde{u}_t^i)]d\psi + \tilde{\lambda} C_t \leq (1 - \delta)\tilde{K}_t + F(\tilde{K}_t, N_t) \quad t = 0, 1, \dots \quad (20)$$

and

$$\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, 0).$$

Let $\tilde{C}_t \equiv \int \mathbb{E}^i[c(\tilde{u}_t^i)]d\psi$ denote aggregate consumption under the optimized allocation. The optimal allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ in this program is Δ -efficient and saves an amount $\tilde{\lambda} C_t$ of aggregate resources in every period. Our measure of the distance of the baseline allocation from the Δ -efficient frontier is $\tilde{\lambda}$.

4.2 Relaxed Problem

We now construct a relaxed planning problem that replaces the resource constraints with a single present value condition. This problem is indexed by a sequence of intertemporal prices or interest rates, which encodes the scarcity of aggregate resources in every period. The relaxed planning problem can be further decomposed in a series of component planning problems corresponding to the different types $i \in I$.

Given some positive intertemporal prices $\{\tilde{Q}_t\}$ with the normalization that $\sum_{t=0} \tilde{Q}_t C_t = 1$ we replace the sequence of resource constraints (20) with the single present value con-

dition

$$\tilde{\lambda} = \sum_{t=0}^{\infty} \tilde{Q}_t \left(F(\tilde{K}_t, N_t) + (1 - \delta)\tilde{K}_t - \tilde{K}_{t+1} - \int \mathbb{E}^i [c(\tilde{u}_t^i)] d\psi \right) \quad (21)$$

which is obtained by multiplying equation (20) by \tilde{Q}_t and summing over $t = 0, 1, \dots$. Formally, the *relaxed planning problem* seeks the allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ where $\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, 0)$ that maximizes $\tilde{\lambda}$ given by equation (21).

The connection with the original planning problem is the following. Suppose that, given some $\{\tilde{Q}_t\}$, the optimal allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ for the relaxed problem satisfies the resource constraints (20). Then, this allocation solves the original planning problem. This relaxed problem approach is adapted from Farhi and Werning [2007] and is related to the first welfare theorem proved in Atkeson and Lucas [1992].

The converse is also true. Indeed, the prices $\{\tilde{Q}_t\}$ are Lagrange multipliers and Lagrangian necessity theorems guarantee the existence of prices $\{\tilde{Q}_t\}$ for the relaxed problem. The following lemma, which follows immediately from Theorem 1, Section 8.3 in Luenberger [1969], provides one such result.

Lemma 1. *Suppose $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ solves the planning problem and the resource constraints (20) hold with equality. Then there exists a sequence of prices $\{\tilde{Q}_t\}$ such that this same allocation solves the relaxed planning problem.*

The relaxed planning problem can be decomposed into a subproblem for capital $\{\tilde{K}_t\}$ and a series of component planning problems for the utility assignment $\{\tilde{u}_t^i\}$. The subproblem for capital maximizes the right hand side of equation (21) with respect to $\{\tilde{K}_{t+1}\}$. The first-order conditions, which are necessary and sufficient for an interior optimum, are

$$1 = \tilde{q}_t (F_K(\tilde{K}_{t+1}, N_{t+1}) + 1 - \delta) \quad t = 0, 1, \dots$$

where $\tilde{q}_t \equiv \tilde{Q}_{t+1}/\tilde{Q}_t$. This, together with the normalization that $\sum_{t=0}^{\infty} \tilde{Q}_t C_t = 1$, implies a one-to-one relationship between $\{\tilde{K}_{t+1}\}$ and $\{\tilde{Q}_t\}$.

The component planning problem for type $i \in I$ maximizes the right hand side of equation (21) with respect to the utility assignment $\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, 0)$. The objective reduces to minimizing the present value of consumption:

$$\sum_{t=0}^{\infty} \frac{\tilde{Q}_t}{\tilde{Q}_0} \mathbb{E}^i [c(\tilde{u}_t^i)] \quad (22)$$

Each of these component planning problems can be solved independently. Since both the objective and the constraints are convex it follows immediately that, given $\{\tilde{q}_t\}$, the first

order conditions for optimality at an interior solution in the component planning problem (22) coincide exactly with the Inverse Euler equation (18):

$$c'(\tilde{u}_t^i) = \frac{\tilde{q}_t}{\beta} \mathbb{E}_t^i[c'(\tilde{u}_{t+1}^i)] \quad t = 0, 1, \dots$$

4.3 A Bellman Equation

In most situations of interest, the baseline allocation admits a recursive representation for some endogenous state variable. This is the case whenever $\{\theta_t^i\}$ is a Markov process and the baseline allocation depends on the history of shocks $\theta^{i,t-1}$ in a way that can be summarized by an endogenous state x_t^i , with law of motion $x_t^i = M(x_{t-1}^i, \theta_t^i)$ and given initial condition x_0^i . Note that the transition matrix M is assumed to be independent of i . Types then correspond to different initial value x_0^i . This is the only thing that distinguishes them. The endogenous state x_t^i is a function of the history of exogenous shocks $\theta^{i,t}$. Defining the state vector $s_t^i = (x_t^i, \theta_t^i)$ there must exist a function \bar{u} such that $u_t^i(\theta^{i,t}) = \bar{u}(s_t^i)$ for all $\theta^{i,t}$. In what follows we drop the hat notation and we stop indexing the allocation by the type i of the agents. We denote a baseline allocation by $u(s_t)$ and use the notation $c(s_t)$ for $c(u(s_t))$.

The requirement that the baseline allocation be recursive in this way is hardly restrictive. Of course, the endogenous state and its law of motion depend on the particular economic model generating the baseline allocation. A leading example in this paper is the case of incomplete markets Bewley economies in Huggett [1993] and Aiyagari [1994]. In these models, described in more detail in Section 6, each individual is subject to an exogenous Markov process for income or productivity and saves using a riskless asset. At a steady state, the interest rate on this asset is constant, so that the agent's solution can be summarized by a stationary savings rule. The baseline allocation can then be summarized using asset wealth as an endogenous state, with law of motion M given by the agent's optimal saving rule.¹³

For such baseline allocations we can reformulate the component planning problems recursively as follows.

The idea is to take s_t as an exogenous process and keep track of the additional lifetime utility Δ_{t-1} previously promised as an endogeneous state variable.

For any date τ define the continuation plans $\{u_t^\tau\}_{t=\tau}^\infty$ with $u_t^\tau(s) = u_{t+\tau}(s)$ and the value function corresponding to the component planning problem starting at date τ with

¹³ Another example are allocations generated by a dynamic contract. The state variable then includes the promised continuation utility [see Spear and Srivastava, 1987] along with the exogenous state.

state s

$$K(\Delta_-, s; \tau) \equiv \inf_{\Delta} \sum_{t=0}^{\infty} \frac{\bar{Q}_{\tau+t}}{\bar{Q}_{\tau}} \mathbb{E}[c(\tilde{u}_t^{\tau}) \mid s_{\tau} = s] \quad \text{s.t. } \{\tilde{u}_t^{\tau}\} \in Y(\{u_t^{\tau}\}, \Delta_-).$$

This value function satisfies the Bellman equation

$$K(\Delta_-, s; \tau) = \min_{\Delta} [c(u(s) + \Delta_- - \beta\Delta) + \bar{q}_{\tau} \mathbb{E}[K(\Delta, s'; \tau + 1) \mid s]]. \quad (23)$$

The optimization over Δ is one-dimensional and convex and delivers a policy function $g(\Delta_-, s; \tau)$. Combining the necessary and sufficient first-order condition for Δ with the envelope condition yields the Inverse Euler equation in recursive form

$$c'(u(s) + \Delta_- - \beta g(\Delta_-, s; \tau)) = \frac{\bar{q}_{\tau}}{\beta} \mathbb{E}[c'(u(s') + g(\Delta_-, s; \tau) - \beta g(g(\Delta_-, s; \tau), s'; \tau + 1)) \mid s]$$

This condition can be used to compute $g(\cdot, \cdot; \tau)$ for given $g(\cdot, \cdot; \tau + 1)$.

An optimal plan for $\{\tilde{u}_t\}$ can then be generated from the sequence of policy functions $\{g(\cdot, \cdot; \tau)\}$ by setting $\tilde{u}(s^t) = u(s_t) + \Delta(s^{t-1}) - \beta\Delta(s^t)$ and using the recursion $\Delta(s^t) = g(\Delta(s^{t-1}), s_t; t)$ with initial condition $\Delta_{-1} = 0$.

With a fixed discount factor q , the value function is independent of time, $K(\Delta_-, s)$, and solves the stationary Bellman equation

$$K(\Delta_-, s) = \min_{\Delta} [c(u(s) + \Delta_- - \beta\Delta) + q \mathbb{E}[K(\Delta, s') \mid s]]. \quad (24)$$

This dynamic program admits an analogy with a consumer's income fluctuation problem that is both convenient and enlightening. We transform variables by changing signs and switch the minimization to a maximization. Let $\tilde{\Delta}_- \equiv -\Delta_-$, $\tilde{K}(\tilde{\Delta}_-; s) \equiv -K(-\tilde{\Delta}_-; s)$ and $\tilde{U}(x) \equiv -c(-x)$. Note that the pseudo utility function \tilde{U} is increasing, concave and satisfies Inada conditions at the extremes of its domain.¹⁴ Reexpressing the Bellman equation (24) using these transformation yields:

$$\tilde{K}(\tilde{\Delta}_-; s) = \max_{\tilde{\Delta}} [\tilde{U}(-u(s) + \tilde{\Delta}_- - \beta\tilde{\Delta}) + q \mathbb{E}[\tilde{K}(\tilde{\Delta}; s') \mid s]].$$

This reformulation can be read as the problem of a consumer with a constant discount

¹⁴ An important case is when the original utility function is CRRA $U(c) = c^{1-\sigma}/(1-\sigma)$ for $\sigma > 0$ and $c \geq 0$. Then for $\sigma > 1$ the function $\tilde{U}(x)$ is proportional to a CRRA with coefficient of relative risk aversion $\tilde{\sigma} = \sigma/(\sigma - 1)$ and $x \in (0, \infty)$. For $\sigma < 1$ the pseudo utility \tilde{U} is "quadratic-like", in that it is proportional to $-(-x)^{\rho}$ for some $\rho > 1$, and $x \in (-\infty, 0]$.

factor q facing a constant gross interest rate $1 + r = \beta^{-1}$, entering the period with pseudo financial wealth $\tilde{\Delta}_-$, receives a pseudo labor income shock $-u(s)$. The fictitious consumer must decide how much to save $\beta\tilde{\Delta}$; pseudo consumption is then $x = -u(s) + \tilde{\Delta}_- - \beta\tilde{\Delta}$.

The benefit of this analogy is that the income fluctuations problem has been extensively studied and used; it is at the heart of most general equilibrium incomplete market models [Aiyagari, 1994].

With logarithmic utility, the Bellman equation can be simplified considerably. The idea is best seen through the analogy, noting that the pseudo utility function is exponential $-e^{-x}$ in this case. It is well known that for a consumer with CARA preferences a one unit increase in financial wealth, $\tilde{\Delta}$, results in an increase in pseudo-consumption, x , of $r/(1+r) = 1 - \beta$ in parallel across all periods and states of nature. It is not hard to see that this implies that the value function takes the form $\tilde{K}(\tilde{\Delta}_-; s) = e^{-(1-\beta)\tilde{\Delta}_-} \tilde{k}(s)$. These ideas are behind the following result.

Proposition 3. *With logarithmic utility and constant discount q , the value function in equation (24) is given by*

$$K(\Delta_-; s) = e^{(1-\beta)\Delta_-} k(s),$$

where function $k(s)$ solves the Bellman equation

$$k(s) = A c(s)^{1-\beta} (\mathbb{E}[k(s') \mid s])^\beta, \quad (25)$$

where $A \equiv (q/\beta)^\beta / (1-\beta)^{1-\beta}$.

The optimal policy for Δ can be obtained from $k(s)$ using

$$\Delta = \Delta_- - \frac{1}{\beta} \log \left(\frac{(1-\beta)k(s)}{c(s)} \right). \quad (26)$$

Proof. See Appendix B. ■

This solution is nearly closed form: one needs only compute $k(s)$ using the recursion in equation (25), which requires no optimization. No simplifications on the stochastic process for skills are required.

4.4 Idiosyncratic Planning Problem

The full planning problem maximizes over utility assignments and capital. It is useful to also consider a version of the problem that takes the baseline sequence of capital $\{K_t\}$ as

given. Thus, define the *idiosyncratic planning problem* as maximizing λ^I subject to

$$\int \mathbb{E}^i[c(\hat{u}_t^i)]d\psi + \bar{\lambda}^I C_t = C_t \quad t = 0, 1, \dots$$

and $\{\hat{u}_t^i\} \in Y(\{u_t^i\}, 0)$. The efficiency gains λ^I represent the constant proportional reduction in consumption that is possible without changing the aggregate sequence of capital. Of course, the total efficiency gains are larger than the idiosyncratic ones: $\bar{\lambda} \geq \bar{\lambda}^I$.

The idiosyncratic planning problem lends itself to a similar analysis. In order to avoid repetitions, we only sketch the corresponding analysis. We can define a corresponding relaxed problem, given a sequence of prices $\{\hat{Q}_t\}$. We can also prove an analogue of Lemma (1). This relaxed planning problem can then be decomposed into a series of component planning problems. When the baseline allocation is recursive, and given a sequence of prices $\{\hat{Q}_t\}$, we can study the component planning problems using a Bellman equation as in equation (23). The corresponding first-order conditions also take the form of an Inverse Euler equation

$$c'(\hat{u}_t^i) = \frac{\hat{q}_t}{\beta} \mathbb{E}_t^i[c'(\hat{u}_{t+1}^i)] \quad t = 0, 1, \dots$$

where $\hat{q}_t \equiv \hat{Q}_{t+1}/\hat{Q}_t$. In other words, the marginal rates of substitution corresponding to the Inverse Euler equation $\mathbb{E}_t^i[c'(\hat{u}_{t+1}^i) / (\beta c'(\hat{u}_t^i))]$ must be equalized across types and histories in every period. These marginal rates of substitution, however, are not necessarily linked to any technological rate of transformation as in equation (19). The idiosyncratic efficiency gains thus correspond to the gains from equalizing the marginal rate of substitution across types and histories in every period, without changing the sequence of capital.

5 Idiosyncratic and Aggregate Gains with Log Utility

In this section, we focus on the case of logarithmic utility. We first show that the planning problem can be decomposed into an Idiosyncratic planning problem and a simple Aggregate planning problem. We then illustrate this result in the simple benchmark case where the baseline allocation of consumption is a geometric random walk.

5.1 A Decomposition: Idiosyncratic and Aggregate

When utility is logarithmic, our parallel shifts in utility imply proportional shifts in consumption. This allows us to prove a decomposition result: the efficiency gains $\tilde{\lambda}$ using a simple Aggregate planning problem involving the idiosyncratic efficiency gains $\tilde{\lambda}^I$. It also makes it possible to solve the idiosyncratic efficiency gains and the corresponding allocation in closed form when the baseline allocation is recursive and features constant consumption. In this case, the optimum in the planning problem and the efficiency gains $\tilde{\lambda}$ can be solved out almost explicitly by combining the solution of the Idiosyncratic problem with that of the Aggregate planning problem.

Given $\tilde{\lambda}^I \in [0, 1)$, the Aggregate planning problem seeks to determine the aggregate allocation $\{\tilde{C}_t, N_t, \tilde{K}_t, \tilde{\lambda}C_t\}$ that maximizes $\tilde{\lambda}$, subject to $\tilde{K}_0 = K_0$,

$$\tilde{K}_{t+1} + \tilde{C}_t + \tilde{\lambda}C_t \leq (1 - \delta)\tilde{K}_t + F(\tilde{K}_t, N_t) \quad t = 0, 1, \dots$$

and

$$\sum_{t=0}^{\infty} \beta^t U(\tilde{C}_t) = \sum_{t=0}^{\infty} \beta^t U(C_t(1 - \tilde{\lambda}^I))$$

Proposition 4. *Consider a baseline allocation $\{u_t^i, n_t^i, K_t, 0\}$ such that $\sum_{t=0}^{\infty} \beta^t U(C_t)$ is well defined and finite. Suppose that the optimum in the Idiosyncratic planning problem are such that the resource constraints hold with equality. Let $\tilde{\lambda}^I$ denote the idiosyncratic efficiency gains identified by the Idiosyncratic problem. Then the total efficiency gains $\tilde{\lambda}$ underlying the planning problem are determined by the Aggregate planning problem.*

Proof. See Appendix C. ■

The proof of Proposition 4 also establishes that the utility assignment $\{\hat{u}_t^i\}$ that solves the Idiosyncratic planning problem and the utility assignment $\{\tilde{u}_t^i\}$ that solves the original planning problem are related by $\tilde{u}_t^i = \hat{u}_t^i + \delta_t$ where $\delta_t = U(\tilde{C}_t) - U(C_t(1 - \tilde{\lambda}^I))$.

The efficiency gains $\tilde{\lambda}$ can then be decomposed into idiosyncratic efficiency gains $\tilde{\lambda}^I$ and aggregate efficiency gains $\tilde{\lambda}^A$. Aggregate efficiency gains are simply defined as $\tilde{\lambda}^A = \tilde{\lambda} - \tilde{\lambda}^I$.

The analysis of the evolution of the aggregate allocation $\{\tilde{C}_t, \tilde{K}_t\}$ only requires the knowledge of $\tilde{\lambda}^I$, but can otherwise be conducted separately from the analysis of the idiosyncratic problem. The aggregate planning problem is simply that of a standard deterministic growth model, which, needless to say, is straightforward to solve. For example, suppose that the baseline allocation represents a steady state with constant aggregates, $C_t = C_{ss}$, $K_t = K_{ss}$ and $N_t = N_{ss}$. Then the optimized allocation aggregate allocation

$\{\tilde{C}_t, \tilde{K}_t\}$ converges to a steady state $\{\tilde{C}_{ss}, \tilde{K}_{ss}\}$ such that $1 - \delta + F_K(\tilde{K}_{ss}, N_{ss}) = 1/\beta$ and $\tilde{C}_{ss} = F(\tilde{K}_{ss}) - \delta\tilde{K}_{ss} - \tilde{\lambda}C_{ss}$. We will put that result to use in Section 6.

Suppose that the baseline allocation is recursive with state s and features constant aggregate consumption $C_t = C_{ss}$. It is not necessary for the argument that aggregate capital K_t and labor N_t be constant at the baseline allocation. The idiosyncratic efficiency gains can then be easily computed. Proposition 2 implies that we can use $\hat{q}_t = \beta$ to compute the idiosyncratic allocation. Proposition 3 then allows us to compute the solution in closed form. This leads to

$$\tilde{\lambda}^I = 1 - \frac{(1 - \beta) \int k(s) d\psi}{C_{ss}}$$

where $k(s)$ solves equation (24) with $q = \beta$. We can then apply Proposition 4 and compute the efficiency gains $\tilde{\lambda}$ and the corresponding Δ -efficient allocation by solving the Aggregate planning problem, a version of the neoclassical growth model. This provides a complete solution to the planning problem when the baseline allocation is recursive and features constant aggregate consumption.

5.2 Example: Steady States with Geometric Random Walk

Although the main virtue of our approach is that we can flexibly apply it to various baseline allocations, in this section we begin with a simple and instructive case. We maintain the assumption of logarithmic utility throughout. We take the baseline allocation to be a geometric random walk: $s_{t+1} = \varepsilon_t s_t$ with ε_t i.i.d. and $c(s) = s$, so that $u(s) = U(s)$. Moreover, we assume that $\log(\varepsilon)$ is normally distributed with variance σ_ε^2 so that $\mathbb{E}[\varepsilon] \cdot \mathbb{E}[\varepsilon^{-1}] = \exp(\sigma_\varepsilon^2)$. We also assume that the baseline allocation represents a steady state with constant aggregates $C_t = C_{ss}$, $K_t = K_{ss}$ and $N_t = N_{ss}$ which require $\mathbb{E}[\varepsilon] = 1$. We define $r_{ss} = F_K(K_{ss}, N_{ss}) - \delta$ and $q_{ss} = 1/(1 + r_{ss})$. Moreover, we assume that the Euler equation holds at the baseline allocation, which requires $q_{ss} = \beta \mathbb{E}[\varepsilon^{-1}] = \beta \exp(\sigma_\varepsilon^2)$.

Although extremely stylized, a random walk is an important conceptual and empirical benchmark. First, most theories—starting with the simplest permanent income hypothesis—predict that consumption should be close to a random walk. Second, some authors have argued that the empirical evidence on income, which is a major determinant for consumption, and consumption itself shows the importance of a highly persistent component [e.g. Storesletten, Telmer, and Yaron, 2004a]. For these reasons, a parsimonious statistical specification for consumption may favor a random walk.

The advantage is that we obtain closed-form solutions for the optimized allocation, the intertemporal wedge and the efficiency gains. The transparency of the exercise reveals important determinants for the magnitude of efficiency gains. A geometric random

walk however, is special for the following reason. If we apply the decomposition of Section 5.1, then the idiosyncratic efficiency gains are zero. The entirety of the efficiency gains are aggregate. This is because at the baseline allocation, the marginal rates of substitution $\mathbb{E}_t^i [c'(u_{t+1}^i) / (\beta c'(u_t^i))]$ are already equalized across types and histories to β^{-1} . Therefore, the sequence of prices $\{\hat{Q}_t\}$ given by $\hat{Q}_t = \hat{Q}_0 \beta^t$ and $\hat{Q}_0 = (1 - \beta) / C_{ss}$ solve the relaxed version of the Idiosyncratic problem. These marginal rates of substitution are not equalized, however, to the marginal rate of transformation $1 - \delta + F_K(K_{ss}, N_{ss})$. Therefore, this section can be seen as an exploration of the determinants of aggregate efficiency gains.

An Example Economy. Indeed, one can construct an example economy where a geometric random walk for consumption arises as a competitive equilibrium with incomplete markets. To see this, suppose that individuals have logarithmic utility over consumption and disutility from labor $V(n; \theta) = v(n/\theta)$ for some convex function $v(n)$ over work effort n , so that θ can be interpreted as productivity. Skills evolve as a geometric random walk, so that $\theta_{t+1} = \varepsilon_{t+1} \theta_t$, where ε_{t+1} is i.i.d. Individuals can only accumulate a riskless asset paying return q^{-1} equal to the rate of return on the economy's linear savings technology. They face the sequence of budget constraints

$$a_{t+1} + c_t \leq q^{-1} a_t + \theta_t n_t \quad t = 0, 1, \dots$$

and the borrowing constraint that $a_t \geq 0$. Suppose the rate of return q^{-1} is such that $1 \geq \beta q^{-1} \mathbb{E}[\varepsilon^{-1}]$. Finally, suppose that individuals have no initial assets, so that $a_0 = 0$.

In the competitive equilibrium of this example individuals exert a constant work effort n_{ss} , satisfying $n_{ss} v'(n_{ss}) = 1$, and consume all their labor income each period, $c_t = \theta_t n_{ss}$; no assets are accumulated, a_t remains at zero. This follows because the construction ensured that the agent's intertemporal Euler equation holds at the proposed equilibrium consumption process. The level of work effort n_{ss} is defined so that the intra-period consumption-leisure optimality condition holds. Then, since the agent's problem is convex, it follows that this allocation is optimal for individuals. Since the resource constraint trivially holds, it is an equilibrium.

Although this is certainly a very special example economy, it illustrates that a geometric random walk for consumption is a possible equilibrium outcome.

Partial Equilibrium: Linear Technology. We first study the case where the technology is linear with a rate of return $q^{-1} > 1$. This imposes $q = q_{ss} = \beta \exp(\sigma_\varepsilon^2)$. Note that $q < 1$ imposes, for a given discount rate β , an upper bound on the variance of the shocks $\exp(\sigma_\varepsilon^2) < \beta^{-1}$.

Since the idiosyncratic efficiency gains $\tilde{\lambda}^I$ are zero, the solution of the planning prob-

lem can be derived by studying the Aggregate planning problem, which takes a remarkably simple form. The aggregate consumption sequence $\{\tilde{C}_t\}$ that solves the Aggregate planning problem is given by

$$\tilde{C}_t = C_{ss} \exp\left(\frac{\beta}{1-\beta}\sigma_\varepsilon^2\right) \exp(-t\sigma_\varepsilon^2).$$

The efficiency gains $\tilde{\lambda}$ and the optimal utility assignment $\{\tilde{u}\}$ are then readily computed. We can also derive the intertemporal wedge τ that measures the savings distortions at the optimal allocation $U'(c(\tilde{u}(s^t))) = \beta(1-\tau)q^{-1}\mathbb{E}[U'(c(\tilde{u}(s^{t+1})))|s^t]$.

Proposition 5. *Suppose that utility is logarithmic and that the technology is linear. Suppose that the baseline allocation is a geometric random walk with constant aggregate consumption $C_t = C_{ss}$, that the shocks ε are lognormal with variance σ_ε^2 and that the Euler equation holds at the baseline. Then the consumption assignment of the solution of the planning problem is given by $\tilde{c}(s^t) = \exp\left(\frac{\beta}{1-\beta}\sigma_\varepsilon^2\right) \exp(-t\sigma_\varepsilon^2) s_t$. The intertemporal wedge at the optimal allocation is given by $\tau = 1 - \exp(-\sigma_\varepsilon^2)$. The efficiency gains are given by $\tilde{\lambda} = 1 - \frac{\beta^{-1} - \exp \sigma_\varepsilon^2}{\beta^{-1} - 1} \exp\left(\frac{\beta}{1-\beta}\sigma_\varepsilon^2\right)$.*

The optimized allocation has a lower drift than the baseline allocation. Intuitively, our perturbations based on parallel shifts in utility can be understood as allowing consumers to borrow and save with an artificial idiosyncratic asset, the payoff of which is correlated with their baseline idiosyncratic consumption process: they can increase their consumption today by reducing their consumption tomorrow in such a way that they reduce their consumption tomorrow more in states where consumption is high than in states where consumption is low. The desirable insurance properties of these perturbation make them attractive, leading to a front-loading of consumption. In other words, because our perturbations allow for better insurance, they reduce the benefits of engaging in precautionary savings by accumulating a buffer stock of risk free assets. As a result, it is optimal to front-load consumption, by superimposing a downward drift $\exp(-\sigma_\varepsilon^2)$ on the baseline allocation, where the variance in the growth rate of consumption σ_ε^2 indexes the strength of the precautionary savings motive at the baseline allocation.

In this example, the Inverse Euler equation provides a rationale for a constant and positive wedge $\tau = 1 - \exp(-\sigma_\varepsilon^2)$ in the agent's Euler equation. This is in stark contrast to the Chamley-Judd benchmark result, where no such distortion is optimal in the long run, so that agents are allowed to save freely at the social rate of return.

The efficiency gains are increasing in σ_ε^2 . Note that when $\sigma_\varepsilon^2 = 0$, there are no efficiency gains. For small values of σ_ε^2 , the wedge is given by $\tau \approx \sigma_\varepsilon^2$. The formula for the efficiency gains then takes the form of a simple Ramsey formula $\tilde{\lambda} \approx (\beta/(1-\beta)^2) \tau^2/2$. At the

other extreme, as $\sigma_\varepsilon^2 \rightarrow -\log(\beta)$ the efficiency gains asymptote 100%. The reason is that then $q \rightarrow 1$, implying that the present value of the baseline consumption allocation goes to infinity; in contrast, the cost of the optimal allocation remains finite.

General Equilibrium: Concave Technologies. In this section, we maintain the assumption that utility is logarithmic. We also assume that the baseline allocation is a geometric random walk representing a steady state with constant aggregates, $C_t = C_{ss}$, $K_t = K_{ss}$ and $N_t = N_{ss}$, that the shocks ε are lognormally distributed and that the Euler equation holds at the baseline allocation. We depart from the partial equilibrium assumption of a linear technology, and consider instead the case of concave accumulation technologies. We argue that the efficiency effects may be greatly reduced. This point is certainly not surprising, nor is it specific to the model or forces emphasized here. Indeed, a similar issue arises in the Ramsey literature, the quantitative effects of taxing capital greatly depend on the underlying technology.¹⁵ Unsurprising as it may, it is important to confront this issue to reach meaningful quantitative conclusions.¹⁶

The point that general equilibrium considerations are important can be made most clearly from the following example. We consider the extreme case of a constant endowment: the economy has no savings technology, so that $C_t \leq N_t$ for $t = 0, 1, \dots$ [Huggett, 1993]. Then the baseline allocation is Δ -efficient. This follows since one finds a sequence of intertemporal prices \tilde{q}_t such that the Inverse Euler equation (18) holds. Thus, in this exchange economy there are no efficiency gains from changing the allocation.¹⁷ Certainly the fixed endowment case is an extreme example, but it serves to illustrate that general equilibrium considerations are extremely important.

Consider now a neoclassical production function $F(K, N)$. Applying the results in Section 5.1, we can decompose the planning problem into a idiosyncratic planning problem and an aggregate planning problem, with a corresponding decomposition for efficiency gains. We have already argued that the corresponding efficiency gains are equal to zero $\lambda^I = 0$. Therefore, and just as in the partial equilibrium case, all the efficiency gains are aggregate.

The aggregate planning problem is extremely simple. It seeks to determine the ag-

¹⁵ Indeed, Stokey and Rebelo [1995] discuss the effects of capital taxation in representative agent endogenous growth models. They show that the effects on growth depend critically on a number of model specifications. They then argue in favor of specifications with very small growth effects, suggesting that a neoclassical growth model with exogenous growth may provide an accurate approximation.

¹⁶ A similar point is at the heart of Aiyagari's [1994] paper, which quantified the effects on aggregate savings of uncertainty with incomplete markets. He showed that for given interest rates the effects could be enormous, but that the effects were relatively moderate in the resulting equilibrium of the neoclassical growth model.

¹⁷ This point holds more generally in an endowment economy with CRRA utility when the baseline allocation is a geometric random walk.

gregate allocation $\{\tilde{C}_t, N_t, \tilde{K}_t, \tilde{\lambda}C_t\}$ that maximizes the proportional amount of resources saved in every period $\tilde{\lambda}$, subject to the resource constraints

$$\tilde{K}_{t+1} + \tilde{C}_t + \tilde{\lambda}C_{ss} \leq (1 - \delta)\tilde{K}_t + F(\tilde{K}_t, N_{ss}) \quad t = 0, 1, \dots$$

the restriction that the initial capital is equal the initial capital $\tilde{K}_0 = K_0$ of the baseline allocation and the requirement that

$$\sum_{t=0}^{\infty} \beta^t U(\tilde{C}_t) = \frac{U(C_{ss})}{1 - \beta}.$$

This problem is a simple modification of the neoclassical growth model. The solution involves a transition to a steady state with an interest rate equal to $\tilde{r}_{ss} = 1/\beta - 1$. The efficiency gains, which reflect the strength of the precautionary savings motive, depend how much lower than \tilde{r}_{ss} is the interest rate $r_{ss} = 1/(\beta \exp(\sigma_\epsilon^2))$ at the baseline allocation. This, in turn depends on the variance of consumption growth σ_ϵ^2 .

Empirical Evidence. Suppose one wishes to accept the random walk specification of consumption as a useful empirical approximation. What does the available empirical evidence say about the crucial parameter σ_ϵ^2 ?

Unfortunately, the direct empirical evidence on the variance of consumption growth is very scarce, due to the unavailability of good quality panel data for broad categories of consumption.¹⁸ Moreover, much of the variance of consumption growth in panel data may be measurement error or attributable to transitory taste shocks unrelated to the permanent changes we are interested in here.

There are a few papers that, somewhat tangentially, provide some direct evidence on the variance of consumption growth. We briefly review some of this recent work to provide a sense of what is currently available. Using PSID data, Storesletten, Telmer, and Yaron [2004a] find that the variance in the growth rate of the permanent component of food expenditure lies between 1%–4%.¹⁹ Blundell, Pistaferri, and Preston [2004] use PSID data, but impute total consumption from food expenditure. Their estimates imply a variance of consumption growth of around 1%.²⁰ Krueger and Perri [2004] use the panel element in the Consumer Expenditure survey to estimate a statistical model of consumption. At face value, their estimates imply enormous amounts of mobility and a very large

¹⁸ For the United States the PSID provides panel data on food expenditure, and is the most widely used source in studies of consumption requiring panel data. However, recent work by Aguiar and Hurst [2005] show that food expenditure is unlikely to be a good proxy for actual consumption.

¹⁹ See their Table 3, pg. 708.

²⁰ See their footnote 19, pg 21, which refers to the estimated autocovariances from their Table VI.




Figure 2: Efficiency gains in % when baseline consumption is a geometric random walk for consumption and the Euler equation holds as function of σ_ϵ^2 . The left panel shows the gains in partial equilibrium and the right panel shows the gains in general equilibrium. The bottom dotted line corresponds to $\beta = .96$, the middle dashed line to $\beta = .97$ and top plain line to $\beta = .98$.

variance of consumption growth—around 6%–7%—although most of this should be attributed to a transitory, not permanent, component.²¹ In general, these studies reveal the enormous empirical challenges faced in understanding the statistical properties of household consumption dynamics from available panel data.

An interesting indirect source of information is the cohort study by Deaton and Paxson [1994]. This paper finds that the cross-sectional inequality of consumption rises as the cohort ages. The rate of increase then provides indirect evidence for σ_ϵ^2 ; their point estimate implies a value of $\sigma_\epsilon^2 = 0.0069$. However, recent work using a similar methodology finds much lower estimates [Slesnick and Ulker, 2004, Heathcote, Storesletten, and Violante, 2004].

Calibration. We now display the efficiency gains as a function of σ_ϵ^2 . We choose three

²¹ They specify a Markov transition matrix with 9 bins (corresponding to 9 quantiles) for consumption. We thank Fabrizio Perri for providing us with their estimated matrix. Using this matrix we computed that the conditional variance of consumption growth had an average across bins of 0.0646 (this is for the year 2000, the last in their sample; but the results are similar for other years).

possible discount factors $\beta = 0.96, 0.97$ and 0.98 . For the linear technology model, this imposes $q = \beta \exp(\sigma_\varepsilon^2)$. For the neoclassical growth model, we take $F(K, N) = K^\alpha N^{1-\alpha}$ with $\alpha = 1/3$. Figure 2 plots the efficiency gains as a function of $\exp(\sigma_\varepsilon^2)$ for the linear technology model and for the neoclassical growth model. The figure uses an empirically relevant range for $\exp(\sigma_\varepsilon^2)$.

Consider first the linear technology model. For the parameters under consideration, the efficiency gains range from minuscule - less than 0.1% to very large - over 10%. The effect of the discount factor β is nearly equivalent to increasing the variance of shocks; that is, moving from $\beta = .96$ to $\beta = .98$ has the same effect as doubling σ_ε^2 . To understand this, interpret the lower discounting not as a change in the actual subjective discount, but as calibrating the model to a shorter period length. But then holding the variance of the innovation between periods constant implies an increase in uncertainty over any fixed length of time. Clearly, what matters is the amount of uncertainty per unit of discounted time.

Consider now the neoclassical growth model. There again, there is considerable variation in the size of the efficiency gains, depending on the parameters. However, note that the efficiency gains are much smaller than in the linear technology model. Large differences in interest rates $\tilde{r}_{ss} - r_{ss}$ are necessary to generate substantial efficiency gains: it takes a difference of around 2%, to get efficiency gains that are bigger than 1%.

Lessons. Three lessons emerge from our simple exercise. First, efficiency gains are potentially far from trivial. Second, they are quite sensitive to two parameters available in our exercise: the variance in the growth rate of consumption and the subjective discount factor. Third, general equilibrium forces can greatly mitigate them.

We conclude that, in our view, the available empirical evidence does not provide reliable estimates for the variance of the permanent component of consumption growth. Thus, for our purposes, attempts to specify the baseline consumption process directly are impractical. That is, even if one were willing to assume the most stylized and parsimonious statistical specifications for consumption, the problem is that the key parameter remains largely unknown. This suggests that a preferable strategy is to use consumption processes obtained from models that have been successful at matching the available data on consumption and income. It also follows from this discussion that it is crucial to use general equilibrium models is crucial to the quantitative assessment of the magnitude of the efficiency gains.

6 An Aiyagari Economy

We now take the steady state equilibrium allocation of an incomplete market economy. We adopt the calibrations from the seminal work of Aiyagari [1994].

The Aiyagari Economy. Aiyagari [1994] considered a Bewley economy, where a continuum of agents each solve an income fluctuations problem, saving in a risk free asset. Efficiency labor is specified as a first-order autoregressive process in logarithms

$$\log(n_t) = \rho \log(n_{t-1}) + (1 - \rho) \log(n_{ss}) + \varepsilon_t$$

where ε_t is an i.i.d. random variable assumed Normally distributed with mean zero and standard deviation σ_ε . With a continuum of mass one of agents the average efficiency labor supply is $N_{ss} = n_{ss}$.

Labor income is given by the product $w_{ss} n_t$ where w_{ss} is the steady-state wage. Agents face the following sequence of budget constraints

$$a_{t+1} + c_t \leq (1 + r_{ss})a_t + w_{ss}n_t$$

for all $t = 0, 1, \dots$. In addition borrowing is not allowed: $a_t \geq 0$.

The equilibrium steady-state wage is given by the marginal product of labor $w_{ss} = F_N(K_{ss}, N_{ss})$ and the interest rate is given by the net marginal product of capital, $r_{ss} = F_K(K_{ss}, N_{ss}) - \delta$.²² For any interest rate $r_{ss} < \beta^{-1} - 1$, agent optimization leads to an invariant cross-sectional distribution for s_t , which we denote by ψ . A steady-state equilibrium requires average assets, under ψ , to equal the capital stock K_{ss} .

Individual consumption is a function of the state variable $s_t \equiv (a_t, n_t)$ which evolves as a Markov process. We take this as our baseline allocation with agents distinguished by their initial conditions $i = s_0$, distributed according to the invariant distribution ψ .

Numerical Method. To solve the planning problem we use the result developed in Section 4.2. To apply this result, we seek the appropriate sequence of discount factors $\{\tilde{q}_t\}$ as follows. For any given $\{\tilde{q}_t\}$, we solve the non-stationary Bellman equation (23) using a policy iteration method with Δ_{t-1} as the endogenous state variable and s_t as the exogenous state. Using the underlying policy function for consumption in equation (23), and integrating in every period over ψ , we compute an aggregate sequence of consumption $\{\tilde{C}_t\}$. Using the resource constraints, we can solve for $\tilde{\lambda}$ and a sequence for capital $\{\tilde{K}_t\}$

²² For simplicity this assumes no taxation. It is straightforward to introduce taxation. However, we conjecture that since taxation of labor income acts as insurance, it effectively reduces the variance of shocks to net income. Lower uncertainty will then only lower the efficiency gains we compute.

that has $\tilde{K}_0 = K_{ss}$.

From Section 4.2, if the condition $1/\tilde{q}_t = 1 + F_K(\tilde{K}_t, N_{ss}) - \delta$ is met this constitutes a solution. Otherwise, we take a new sequence of discount factors given by $\tilde{q}'_t = (1 + F_K(\tilde{K}_t, N_{ss}) - \delta)^{-1}$ and iterate until convergence to a fixed point.

Calibration. We simulate the economy for all the parameter values considered in Aiyagari [1994]. The discount factor is set to $\beta = .96$, the production function is Cobb-Douglas with a share of capital of 0.36, capital depreciation is 0.08. The utility function is assumed CRRA so that $U(c) = c^{1-\sigma}/(1-\sigma)$ with $\sigma \in \{1, 3, 5\}$. Aiyagari argues, based on various sources of empirical evidence, for a baseline parametrization with a coefficient of autocorrelation of $\rho = 0.6$ and a standard deviation of labor income of 20%. Following Aiyagari, we also consider different values for the coefficient of relative risk aversion, the autocorrelation coefficient $\rho \in \{0, 0.3, 0.6, 0.9\}$ and the standard deviation of log income, $\text{Std}(\log(n_t)) = \sigma_\varepsilon \in \{0.2, 0.4\}$.

These values are based on the following studies. Kydland [1984] finds that the standard deviation σ_ε of annual hours worked from PSID data is around 15%. Using data from the PSID and the NLS, Abowd and Card [1987] and Abowd and Card [1989] find that the standard deviation of percentage changes in real earnings and annual hours are about 40% and 35% respectively. They report a first order serial correlation coefficient ρ of about 0.3, resulting in a estimate of σ_ε of 34%. Using PSID data, Heaton and Lucas [1996] estimate a range of 0.23 to 0.53 for ρ and a range of 27% to 40% for σ_ε .

Some recent studies, within a life cycle context, consider an alternative approach and estimate a process for log earnings indirectly, by using the increase in the observed cross-sectional inequality of earnings over time within a cohort. Two views have been articulated. The first view, dating back to Deaton and Paxson [1994] and developed most recently by Storesletten et al. [2004a] and Storesletten et al. [2004b], posits that the increase in earning inequality over time within a cohort is due to large and persistent income shocks. This leads to estimates of ρ around or above 0.9 and a range of estimates for σ_ε between 0.3 and 0.6, on the high end of the range of parameter values explored by Aiyagari [1994]. The second view, dating back to Lillard and Weiss [1979] and Hause [1980], and developed Guvenen [2007] and Guvenen [2009], argues that the increase in cross sectional earnings inequality over time within a cohort is better explained by an alternative model where agents face individual-specific income profiles. This view leads to lower values for ρ and σ_ε , around 0.8 and 25% respectively. Since both approaches are based on a life-cycle framework, the estimates are not directly relevant for our infinite horizon setup. Indeed, as we have shown, a key object for our analysis is the conditional variance of consumption growth which depends on the conditional variance of permanent income




Figure 3: Path of aggregate consumption for the baseline and the optimized allocations for $\sigma = 1$, $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$ and $\rho = 0.9$.

growth. Life-cycle models incorporate a retirement period, generating a longer horizon for consumption than for labor income. This reduces the impact of permanent income shocks on consumption. Taking this into account, estimates for the persistence ρ and the standard deviation of labor income σ_ϵ derived in the context of a life-cycle model would therefore have to be adjusted downwards in order to be used in our numerical exercises.

Results. We find that the optimized allocation always features aggregates converging to new steady state values: $\tilde{C}_t \rightarrow \tilde{C}_{ss}$, $\tilde{K}_t \rightarrow \tilde{K}_{ss}$, $\tilde{q}_t \rightarrow \tilde{q}_{ss}$ as $t \rightarrow \infty$.²³ Figure 3 plots the path of aggregate consumption for one particular parameter case. Since the baseline allocation represents a steady state, its aggregate consumption is constant. Aggregate consumption for the optimized allocation is initially above this level, but declines monotonically, eventually reaching a new, lower, steady state. For the same parameter values, Figure 4 shows a typical sample path for individual income, cash in hand, consumption and utility. Optimized consumption appears more persistent and displays a downwards trend in the initial periods.

Tables 1, 2 and 3 collect the results of our simulations. All tables report our measure of efficiency gains, $\tilde{\lambda}$. In the logarithmic utility case ($\sigma = 1$), Table 1 includes the idiosyncratic and aggregate components $\tilde{\lambda}^I$ and $\tilde{\lambda}^A$. For references, the tables show the baseline and optimized steady state interest rates r_{ss} and \tilde{r}_{ss} . The second column, showing r_{ss} , replicates Aiyagari's Table II (pg. 678).

²³ In the case of logarithmic utility functions, we provided a formal proof of this result in Section 5.1. In the more general CRRA utility function case, we rely solely on our numerical results.

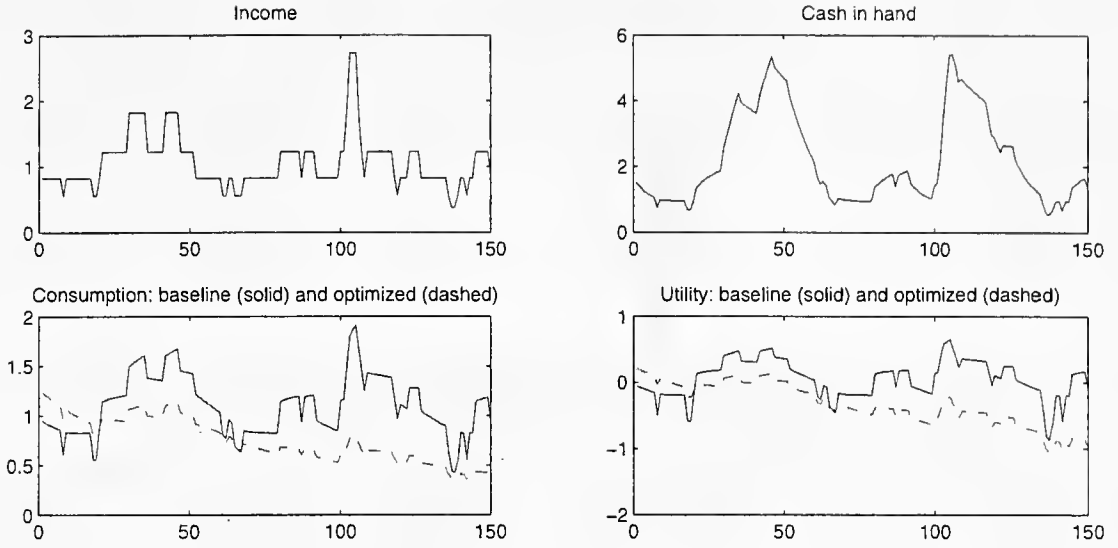


Figure 4: Simulation of a typical individual sample path for $\sigma = 1$, $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$ and $\rho = 0.9$.

The baseline's interest rate r_{ss} is decreasing in the size and persistence of the shocks, as well as in the coefficient of relative risk aversion. Precautionary saving motives are stronger and depress the equilibrium interest rate. The interest rates at the optimized allocation are consistent with the conclusions from Proposition 2. In particular, $\tilde{r}_{ss} = 1/\beta - 1$ when utility is logarithmic, and $\tilde{r}_{ss} > 1/\beta - 1$ when $\sigma > 1$. Moreover, we find that \tilde{r}_{ss} increases with σ and the size and the persistence of the labor income shocks. These comparative statics for \tilde{r}_{ss} are precisely the reverse of those for r_{ss} . This is perhaps not surprising, given the reversal in the sign of the power coefficients in the Inverse Euler and Euler equations, σ and $-\sigma$, respectively.

Efficiency gains are increasing with the variance of the shocks and with their persistence. They also increase with the coefficient of relative risk aversion. Aiyagari argues, based on various sources of empirical evidence, for a parameterization with a coefficient of autocorrelation of $\rho = 0.6$ and a standard deviation of labor income of 20%. For this preferred specification, we find that efficiency gains are small – below 0.2% for all three values for the coefficient of relative risk aversion.

In the logarithmic case, efficiency gains are moderate – always less than 1.3%. Our decomposition along the lines of Section 5.1 shows that idiosyncratic efficiency gains $\tilde{\lambda}^I$ completely dwarf aggregate efficiency gains $\tilde{\lambda}^A$. Our finding that idiosyncratic gains are modest could have perhaps been anticipated by our illustrative geometric random-walk example, where idiosyncratic gains are zero. Intuitively, efficiency gains from the idiosyncratic component require differences in the expected consumption growth rate across in-

$\sigma = 1$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .2$						
Efficiency Gains						
ρ	r_{ss}	\tilde{r}_{ss}	Idiosyncratic	Aggregate	Total	Borrowing
0	4.14%	4.17%	0.0%	0.0%	0.0%	0.0%
0.3	4.13%	4.17%	0.0%	0.0%	0.0%	0.0%
0.6	4.09%	4.17%	0.0%	0.0%	0.0%	0.0%
0.9	3.95%	4.17%	0.2%	0.0%	0.2%	0.1%

$\sigma = 1$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$						
Efficiency Gains						
ρ	r_{ss}	\tilde{r}_{ss}	Idiosyncratic	Aggregate	Total	Borrowing
0	4.06%	4.17%	0.1%	0.0%	0.1%	0.1%
0.3	3.97%	4.17%	0.2%	0.0%	0.2%	0.2%
0.6	3.79%	4.17%	0.4%	0.0%	0.4%	0.3%
0.9	3.38%	4.17%	1.2%	0.1%	1.3%	1.1%

Table 1: Efficiency Gains for replication of Aiyagari [1994] when $\sigma = 1$.

dividuals. When individuals smooth their consumption over time effectively the remaining differences are small—as a result, so are the efficiency gains. Efficiency gains from the aggregate component are directly related to the difference between the equilibrium and optimal steady-state capital. With logarithmic utility this is equivalent, to the difference between the equilibrium steady-state interest rate and $\beta^{-1} - 1$, the interest rate that obtains with complete markets. Hence, our finding of low aggregate efficiency gains is directly related to Aiyagari’s [1994] main conclusion: for shocks that are not implausibly large or for moderate risk aversion, precautionary savings are small in the aggregate, in that steady-state capital and interest rate are close their complete-markets levels, as shown in our Table 1.

For the range of parameters that we consider, the efficiency gains range from minuscule – less than 0.1%– to very large – more than 8.4%. However, efficiency gains larger than 1.3% are only reached for combinations of high values of relative risk aversion – σ greater than 3– and both large and highly persistent shocks – a standard deviation of labor income of 40%, and a mean-reversion coefficient ρ greater than 0.6.

The Role of Borrowing Constraints. The market arrangement in Aiyagari’s economy imposes borrowing constraints that limit the ability to trade consumption intertemporally. At the baseline allocation, the Euler equation holds with equality $U'(c(u_t^i)) = \beta(1 + r_{ss})\mathbb{E}_t^i [U'(c(u_{t+1}^i))]$ holds when assets and current income are high enough. However, for low levels of assets and income, the agent may be borrowing constrained so that

$\sigma = 3$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .2$				
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains	
			Total	Borrowing
0	4.09%	4.21%	0.0%	0.0%
0.3	4.02%	4.37%	0.0%	0.0%
0.6	3.88%	4.45%	0.2%	0.2%
0.9	3.36%	4.75%	0.7%	0.7%

$\sigma = 3$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$				
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains	
			Total	Borrowing
0	3.77%	4.62%	0.2%	0.2%
0.3	3.47%	4.77%	0.5%	0.5%
0.6	2.89%	5.03%	1.2%	1.2%
0.9	1.47%	5.56%	4.2%	3.3%

Table 2: Efficiency Gains for replication of Aiyagari [1994] when $\sigma = 3$.

the Euler conditions holds with strict inequality $U'(c(u_t^i)) > \beta(1 + r_{ss})\mathbb{E}_t^i [U'(c(u_{t+1}^i))]$.

In contrast, since our planning problem does not impose arbitrary restrictions on the perturbations, it places no such limits on intertemporal reallocation of consumption. Thus, the perturbations can effectively undo limits to borrowing. As a result, part of the efficiency gains we compute can be attributed to the relaxation of borrowing constraints, not to the introduction of savings distortions.

To get an idea of the efficiency gains that are obtained from the relaxation of borrowing constraints, we performed the following exercise. The basic idea is to use the perturbations to construct a new allocation where the Euler equation always holds with equality. That is, the new allocation stops short of satisfying the Inverse Euler equation, so it does not introduce positive intertemporal wedges. Instead, it removes the negative intertemporal wedges that were present due to borrowing constraints. We compute the resource savings from this perturbations as a simple measure of the efficiency gains due to the relaxation of borrowing constraints.²⁴

More precisely, we seek to determine the unique allocation $\{\tilde{u}_t^{i,B}, n_t^i, \tilde{K}_t^B, \tilde{\lambda}^B C_{ss}\}$, where B stands for borrowing, that satisfies the following constraints. First, we impose that the allocation be achievable through parallel perturbations of the baseline allocation, and deliver the same utility

$$\{\tilde{u}_t^{i,B}\} \in Y(\{u_t^i\}, 0).$$

²⁴Theoretically, it isn't evident that these perturbations actually produce positive efficiency gains, as opposed to negative ones. In our cases, the measure always came out to be positive.

$\sigma = 5$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .2$				
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains	
			Total	Borrowing
0	4.01%	5.34%	0.0%	0.0%
0.3	3.89%	5.39%	0.0%	0.0%
0.6	3.61%	5.48%	0.2%	0.2%
0.9	2.66%	5.72%	1.0%	1.0%

$\sigma = 5$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$				
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains	
			Total	Borrowing
0	3.43%	5.54%	0.2%	0.2%
0.3	2.90%	5.55%	0.7%	0.7%
0.6	1.95%	5.58%	2.1%	2.0%
0.9	-0.16%	5.52%	8.4%	6.3%

Table 3: Efficiency Gains for replication of Aiyagari [1994] when $\sigma = 5$.

Second, we impose that the resource constraints hold

$$\tilde{K}_{t+1}^B + \int \mathbb{E}^i \left[c(\tilde{u}_t^{i,B}) \Pr(\theta^{i,t}) \right] d\psi + \tilde{\lambda}^B C_{ss} \leq (1 - \delta) \tilde{K}_t^B + F(\tilde{K}_t^B, N_{ss}) \quad t = 0, 1, \dots$$

and require that the initial capital be equal the initial capital $\tilde{K}_0^B = K_0$ of the baseline allocation. Finally, we impose that the Euler equation hold for every agent in every period

$$U'(c(\tilde{u}_t^{i,B})) = \beta[1 - \delta + F_K(\tilde{K}_{t+1}^B, N_{ss})] \mathbb{E}_t^i[U'(c(u_{t+1}^{i,B}))].$$

The allocation $\{\tilde{u}_t^{i,B}, n_t^i, \tilde{K}_t^B, \tilde{\lambda}^B C_t\}$ can improve on the baseline allocation by insuring that the Euler equation holds for every agent in every period. Note that this allocation does not satisfy the Inverse Euler equation and is therefore not Δ -efficient. We adopt $\tilde{\lambda}^B$ as our measure of the efficiency gains deriving from the relaxation of borrowing constraints.

The rightmost column in each of table reports the efficiency gains $\tilde{\lambda}^B$ that can be attributed to the relaxation of borrowing constraints. Perhaps the most important conclusion of our numerical exercise can be drawn by comparing the efficiency gains $\tilde{\lambda}$ with the efficiency gains $\tilde{\lambda}^B$ deriving from the relaxation of borrowing constraints. We find that for all size and persistence of shocks, and all utility specifications, most of the efficiency gains come from the relaxation of borrowing constraints. Indeed, the difference between $\tilde{\lambda}$ and $\tilde{\lambda}^B$ is less than 0.1% except for the largest and most persistent shocks – a standard deviation of labor income growth of 40%, and a mean-reversion coefficient ρ greater than

0.9. In other words, most of the efficiency come from allowing agents to better smooth their consumption over time by alleviating borrowing constraints, rather than by optimally distorting savings as prescribed by the substitution the Inverse Euler equation to the Euler equation.

7 Conclusions

The main contribution of this paper is to provide a method for evaluating the auxiliary role that savings distortions may play in social insurance arrangements. We put it to use to evaluate the welfare importance of recent arguments for distorting savings and capital accumulation based on the Inverse Euler equation.

We believe that the methodological contribution of this paper transcends our own quantitative explorations of it. The method developed here is flexible enough to accommodate several extensions and it may be of interest to investigate how these may affect the quantitative conclusions found here for the benchmark Aiyagari economy. In a separate paper [Farhi and Werning, 2008] we pursued such an extension using Epstein-Zin preferences that separate risk aversion from the intertemporal elasticity of substitution. In ongoing work, we study an environment with overlapping-generations instead of the infinite horizon dynastic setup used here. This introduces intergenerational considerations that are absent in our setting. In particular, with an infinitely-lived agent we have shown that the planner chooses a declining sequence for capital. In an overlapping-generations model this feature, in and of itself, affects future generations negatively. This points towards finding lower potential efficiency gains. However, new opportunities for intertemporal reallocations do arise since consumption can be shifted across generations, for any given aggregate sequence of consumption. Finally, as discussed in Section 6, an overlapping generation requires a different calibration of the income process. Despite these new considerations, the basic methodology and decompositions introduced here are useful for such extensions.

Appendix

A Proofs for Section 3.2

To simplify the arguments, we first assume that the horizon is finite with terminal period $T > 0$. We also assume that $V(n; \theta)$ is continuously differentiable with respect to n and θ . We also assume throughout that shocks are continuous.

Consider an assignment for utility and labor $\{u_t^i, n_t^i\}$. We define the continuation value $U(\theta^{i,t})$ conditional on history $\theta^{i,t}$ as follows

$$U^i(\theta^{i,t}) = \mathbb{E} \left[\sum_{s=0}^T \beta^s \left[u_{t+s}^i(\theta^{i,t+s}) - V(n_{t+s}^i(\theta^{i,t+s}); \theta_{t+s}^i) \right] \middle| \theta^{i,t} \right].$$

Fix a history $\theta^{i,t}$ and consider a report $\hat{\theta}_t^i \in \Theta$. Define the strategy $\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}$ as follows: $\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}(\theta^{i,s}) = \theta^{i,s}$ except if $\theta^{i,s} \succeq \theta^{i,t}$, in which case

$$\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}(\theta^{i,t-1}, \theta_t^i, \theta_{t+1}^i, \theta_{t+2}^i, \dots) = (\theta^{i,t-1}, \hat{\theta}_t^i, \theta_{t+1}^i, \theta_{t+2}^i, \dots).$$

This strategy coincides with truth-telling except in period t after history $\theta^{i,t}$ where the report $\hat{\theta}_t^i$ can be different from the true shock θ_t^i .

For clarity and brevity, we use the notation $\hat{\theta}^{i,s}$ for $\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}(\theta^{i,s})$. We denote the continuation utility after history $\theta^{i,s} \succeq \theta^{i,t}$ under the strategy $\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}$, by $U^i(\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}(\theta^{i,s}))$ or $U^i(\hat{\theta}^{i,s}; \theta^{i,t})$ for short. Similarly, we denote by $\mathbb{E}[U^i(\hat{\theta}^{i,s}; \theta^{i,t})]$ or $\mathbb{E}[U^i(\hat{\theta}^{i,s}) | \theta^{i,t}]$ the expectation of this continuation utility, conditional on the realized history $\theta^{i,t}$ at date t .

We say that the assignment for utility and labor $\{u_t^i, n_t^i\}$ is *regular* if for all $t \geq 0$ and $\theta^{i,t-1} \in \Theta^t$, the continuation utility $U^i(\hat{\theta}^{i,t}; \theta_{i,t})$ is absolutely continuous in the true shock θ_t^i , differentiable with respect to the true shock θ_t^i , and the derivative, which we denote by

$$-\frac{\partial}{\partial \theta_t^i} V(n(\hat{\theta}^{i,t}); \theta_t^i) + \beta \frac{\partial}{\partial \theta_t^i} \mathbb{E}[U^i(\hat{\theta}^{i,t+1}) | \theta^{i,t}]$$

is bounded by a function $b(\theta_t^i; \theta^{i,t-1})$ which is integrable with respect to θ_t^i .

Consider two regular, incentive compatible, assignments for utility and labor $\{u_t^i, n_t^i\}$ and $\{\tilde{u}_t^i, n_t^i\}$, that share the same assignment for labor $\{n_t^i\}$. Denote by $\{U^i\}$ and $\{\tilde{U}^i\}$ the corresponding continuation utilities. We now show that there exists $\Delta_{-1} \in \mathbb{R}$ such that $\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, \Delta_{-1})$.

Note we can restate the result as follows: for all t with $0 \leq t \leq T-1$ and for all

$\theta^{i,t-1} \in \Theta^t$, there exists $\Delta(\theta^{i,t-1}) \in \mathbb{R}$ such that

$$\tilde{U}^i(\theta^{i,t}) = U^i(\theta^{i,t}) + \Delta(\theta^{i,t-1}). \quad (27)$$

Fix a history $\theta^{i,t}$ and consider the strategies $\sigma_{\tilde{\theta}_t^i}^{\theta^{i,t}}$. Using Theorem 4 in Milgrom and Segal [2002], Incentive compatibility and regularity implies the following Envelope condition

$$\begin{aligned} U^i(\theta^{i,t-1}, \theta_t^i) &= U^i(\theta^{i,t-1}, \theta_t^{i'}) \\ &+ \int_{\theta_t^{i'}}^{\theta_t^i} \left[-\frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} V(n(\tilde{\theta}^{i,t}); \tilde{\theta}_t^i) + \beta \frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} \mathbb{E} \left[U^i(\tilde{\theta}^{i,t+1}) | (\tilde{\theta}_{t-1}^i, \tilde{\theta}_t^i) \right] d\tilde{\theta}_t^i \right]. \end{aligned} \quad (28)$$

where $\tilde{\theta}^{i,t+s} = (\theta^{i,t-1}, \tilde{\theta}_t^i, \theta_{t+1}^i, \dots, \theta_{t+s}^i)$ for all $s \geq 0$ and $\tilde{\theta}^{i,t-s} = \theta^{i,t-s}$ for all $s > 0$.

The same expression holds for $\{\tilde{u}_t^i, n_t^i\}$. Using the fact that shocks are continuous and that $\{u_t^i, n_t^i\}$ and $\{\tilde{u}_t^i, n_t^i\}$ share the same labor assignment, this implies that there exists $\Delta(\theta^{i,t-1}) \in \mathbb{R}$ such that

$$\begin{aligned} \tilde{U}^i(\theta^{i,t}) &= U^i(\theta^{i,t}) + \Delta(\theta^{i,t-1}) \\ &+ \beta \int_{\underline{\theta}}^{\theta_t^i} \frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} \mathbb{E} \left[\tilde{U}^i(\tilde{\theta}^{i,t+1}) - U^i(\tilde{\theta}^{i,t+1}) | (\tilde{\theta}_{t-1}^i, \tilde{\theta}_t^i) \right] d\tilde{\theta}_t^i. \end{aligned}$$

For $t = T$, the last term on the right hand side vanishes and we have that for all $\theta^{i,T} \in \Theta^{T+1}$, there exists $\Delta(\theta^{i,T-1}) \in \mathbb{R}$ such that (27) holds. We proceed by induction on t . Suppose that $t \geq 0$ and that for all $\theta^{i,t} \in \Theta^{t+1}$, there exists $\Delta(\theta^{i,t}) \in \mathbb{R}$ such that (27) holds. Then applying (28) and using the fact that

$$\frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} \mathbb{E} \left[\Delta(\tilde{\theta}^{i,t}) | (\tilde{\theta}^{i,t-1}, \tilde{\theta}_t^i) \right] = 0,$$

we immediately find that for all $\theta^{i,t-1} \in \Theta^t$ there exists $\Delta(\theta^{i,t-1}) \in \mathbb{R}$ such that

$$\tilde{U}^i(\theta^{i,t}) = U^i(\theta^{i,t}) + \Delta(\theta^{i,t-1}).$$

The proof follows by induction.

When the horizon is infinite, the regularity conditions must be complemented with a limit condition. For example, expanding our Envelope condition over two periods, we find

$$\tilde{U}^i(\theta^{i,t}) = U^i(\theta^{i,t}) + \Delta(\theta^{i,t-1}) + \beta^2 \int_{\underline{\theta}}^{\theta_t^i} \frac{\partial}{\partial \tilde{\theta}_t^i} |_{\tilde{\theta}_t^i = \tilde{\theta}_t^i}$$

$$\mathbb{E} \left[\int_{\underline{\theta}}^{\tilde{\theta}_{t+1}^i} \frac{\partial}{\partial \tilde{\theta}_{t+1}^i} |_{\tilde{\theta}_{t+1}^i = \overleftarrow{\theta}_{t+1}^i} \mathbb{E} \left[\left(\tilde{U}^i - U^i \right) \left(\overleftarrow{\theta}^{i,t+2} \right) | \left(\overleftarrow{\theta}^{i,t}, \tilde{\theta}^{i,t+1} \right) \right] d \overleftarrow{\theta}_t^i | \left(\tilde{\theta}^{i,t-1}, \tilde{\theta}^{i,t} \right) \right] d \tilde{\theta}_t^i.$$

Consider the sequence which is constructed by iterating our Envelope condition for the utility and labor assignment $\{u_t^i, n_t^i\}$. The first element of the sequence is

$$\beta \int_{\underline{\theta}}^{\theta_t^i} \frac{\partial}{\partial \tilde{\theta}_t^i} |_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} \mathbb{E} \left[U^i(\tilde{\theta}^{i,t+1}) | (\tilde{\theta}^{i,t-1}, \tilde{\theta}^{i,t}) \right] d \tilde{\theta}_t^i.$$

Similarly, the second element of the sequence is

$$\beta^2 \int_{\underline{\theta}}^{\theta_t^i} \frac{\partial}{\partial \tilde{\theta}_t^i} |_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} \mathbb{E} \left[\int_{\underline{\theta}}^{\tilde{\theta}_{t+1}^i} \frac{\partial}{\partial \tilde{\theta}_{t+1}^i} |_{\tilde{\theta}_{t+1}^i = \overleftarrow{\theta}_{t+1}^i} \mathbb{E} \left[U^i(\overleftarrow{\theta}^{i,t+2}) | (\overleftarrow{\theta}^{i,t}, \tilde{\theta}^{i,t+1}) \right] d \overleftarrow{\theta}_t^i | (\tilde{\theta}^{i,t-1}, \tilde{\theta}^{i,t}) \right] d \tilde{\theta}_t^i.$$

Our proof then carries over if the condition that the limit of these terms when the number of iterations goes to infinity is equal to zero is added to the requirements for regularity.

B Proof of Proposition 3

With logarithmic utility the Bellman equation is

$$\begin{aligned} K(s, \Delta_-) &= \min_{\Delta} [s \exp(\Delta_- - \beta \Delta) + q \mathbb{E}[K(s', \Delta) | s]] \\ &= \min_{\Delta} [s \exp((1 - \beta) \Delta_- + \beta(\Delta_- - \Delta)) + q \mathbb{E}[K(s', \Delta) | s]] \end{aligned}$$

Substituting that $K(\Delta_-, s) = k(s) \exp((1 - \beta) \Delta_-)$ gives

$$k(s) \exp((1 - \beta) \Delta_-) = \min_{\Delta} [s \exp((1 - \beta) \Delta_- + \beta(\Delta_- - \Delta)) + q \mathbb{E}[k(s') \exp((1 - \beta) \Delta) | s]],$$

and cancelling terms:

$$\begin{aligned} k(s) &= \min_{\Delta} [s \exp(\beta(\Delta_- - \Delta)) + q \mathbb{E}[k(s') \exp((1 - \beta)(\Delta - \Delta_-)) | s]] \\ &= \min_d [s \exp(-\beta d) + q \mathbb{E}[k(s') \exp((1 - \beta)d) | s]] \end{aligned}$$

$$= \min_d [s \exp(-\beta d) + q \mathbb{E}[k(s') | s] \exp((1 - \beta)d)]$$

where $d \equiv \Delta - \Delta_-$. We can simplify this one dimensional Bellman equation further. Define $\hat{q}(s) \equiv q \mathbb{E}[k(s') | s] / s$ and

$$M(\hat{q}) \equiv \min [\exp(-\beta d) + \hat{q} \exp((1 - \beta)d)].$$

The first-order conditions gives

$$\beta \exp(-\beta d) = \hat{q}(1 - \beta) \exp((1 - \beta)d) \Rightarrow d = \log \frac{\beta}{(1 - \beta)\hat{q}}. \quad (29)$$

Substituting back into the objective we find that

$$\begin{aligned} M(\hat{q}) &= \frac{1}{1 - \beta} \exp(-\beta d) = \frac{1}{1 - \beta} \exp\left(-\beta \log \frac{\beta}{(1 - \beta)\hat{q}}\right) \\ &= \frac{1}{(1 - \beta)^{1 - \beta} \beta^\beta} \hat{q}^\beta = B \hat{q}^\beta, \end{aligned}$$

where B is a constant defined in the obvious way in terms of β .

The operator associated with the Bellman equation is then

$$T[k](s) = s M\left(q \frac{\mathbb{E}[k(s') | s]}{s}\right) = A s^{1 - \beta} (\mathbb{E}[k(s') | s])^\beta,$$

where $A \equiv B q^\beta = (q/\beta)^\beta / (1 - \beta)^{1 - \beta}$.

Combining the Bellman $k(s)/s = M(\hat{q}) = A \hat{q}^\beta$ with equation (29) yields the policy function as a function of $K(s)$. This completes the proof.

C Proof of Proposition 4

Consider the aggregate allocation $\{\tilde{C}_t, N_t, \tilde{K}_t, \tilde{\lambda} C_t\}$ that solves the Aggregate planning problem and the utility assignment $\{\hat{u}_t^i\} \in Y(\{u_t^i\}, 0)$ that solve the Idiosyncratic planning problem as well as the corresponding idiosyncratic efficiency gains $\tilde{\lambda}^I$.

Since the resource constraints hold with equality in the Idiosyncratic planning problem, we know from Lemma 1 that there exists a sequence of prices $\{\hat{Q}_t\}$, such that

$$c'(\hat{u}_t^i) = \frac{\hat{q}_t}{\beta} \mathbb{E}_t^i[c'(\hat{u}_{t+1}^i)] \quad t = 0, 1, \dots \quad (30)$$

where $\hat{q}_t \equiv \hat{Q}_{t+1}/\hat{Q}_t$. Moreover, the sequence $\{\hat{q}_t\}$ is given by $\hat{q}_t = \beta C_t/C_{t+1}$.

The aggregate allocation $\{\tilde{C}_t, N_t, \tilde{K}_t, \tilde{\lambda} C_t\}$ satisfies the necessary and sufficient first-order conditions

$$U'(\tilde{C}_t) = \beta (1 - \delta + F_K(\tilde{K}_{t+1}, N_{t+1})) U'(\tilde{C}_{t+1}) \quad t = 0, 1, \dots \quad (31)$$

Define the following sequence $\{\delta_t\}$

$$\delta_t = -U((1 - \tilde{\lambda}^l)C_t) + U(\tilde{C}_t) \quad t = 0, 1, \dots$$

We have $\sum_{t=0}^{\infty} \beta^t \delta_t = 0$. With this choice of $\{\delta_t\}$, we then define a utility assignment $\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, 0)$ as follows $\tilde{u}_t^i = u_t^i + \delta_t$. The allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ then satisfies all the constraints of the original planning problem. Moreover using $c'(\tilde{u}_t^i) = c'(\delta_t)c'(u_t^i)$, equation (30) and equation (31), we find that

$$c'(\tilde{u}_t^i) = \frac{\tilde{q}_t}{\beta} \mathbb{E}_t^i[c'(\tilde{u}_{t+1}^i)] \quad t = 0, 1, \dots$$

where

$$1 = \tilde{q}_t (F_K(\tilde{K}_{t+1}, N_{t+1}) + 1 - \delta) \quad t = 0, 1, \dots$$

Hence the allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ satisfies the sufficient first order conditions in the planning problem. It therefore represents the optimum.

References

- John M Abowd and David Card. Intertemporal labor supply and long-term employment contracts. *American Economic Review*, 77(1):50–68, March 1987. 36
- John M Abowd and David Card. On the covariance structure of earnings and hours changes. *Econometrica*, 57(2):411–45, March 1989. 36
- Mark Aguiar and Erik Hurst. Consumption versus expenditure. *Journal of Political Economy*, 113(5):919–948, 2005. 32
- S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659–684, 1994. 1, 4, 5, 6, 20, 23, 25, 31, 35, 36, 37, 38, 39, 40, 41, 42
- Stefania Albanesi and Christopher Sleet. Dynamic optimal taxation with private information. *Review of Economic Studies*, 73(1):1–30, 2006. 2
- Andrew Atkeson and Robert E., Jr. Lucas. On efficient distribution with private information. *Review of Economic Studies*, v59(3):427–453, 1992. 4, 22
- Richard Blundell, Luigi Pistaferri, and Ian Preston. Consumption inequality and partial insurance. IFS Working Papers W04/28, Institute for Fiscal Studies, 2004. 32
- Christophe Chamley. Optimal taxation of capital in general equilibrium. *Econometrica*, 54: 607–22, 1986. 2
- Juan Carlos Conesa and Dirk Krueger. On the optimal progressivity of the income tax code. *Journal of Monetary Economics*, forthcoming, 2005. 5
- Julio Davila, Jay H. Hong, Per Krusell, and Jose-Victor Rios-Rull. Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks. PIER Working Paper Archive 05-023, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania, July 2005. available at <http://ideas.repec.org/p/pen/papers/05-023.html>. 6
- Angus Deaton and Christina Paxson. Intertemporal choice and inequality. *Journal of Political Economy*, 102(3):437–67, 1994. 33, 36
- Peter A. Diamond and James A. Mirrlees. A model of social insurance with variable retirement. Working papers 210, Massachusetts Institute of Technology, Department of Economics, 1977. 2, 13
- Emmanuel Farhi and Iván Werning. Inequality and social discounting. *Journal of Political Economy*, 115(3):365–402, 2007. 22
- Emmanuel Farhi and Iván Werning. Optimal savings distortions with recursive preferences. *Journal of Monetary Economics*, 55:21–42, 2008. 42

- John Geanakoplos and Heracles M. Polemarchakis. Existence, regularity, and constrained suboptimality of competitive allocations when the asset market is incomplete. Cowles Foundation Discussion Papers 764, Cowles Foundation, Yale University, 1985. available at <http://ideas.repec.org/p/cwl/cwldpp/764.html>. 5
- Mikhail Golosov and Aleh Tsyvinski. Designing optimal disability insurance: A case for asset testing. *Journal of Political Economy*, 114(2):257–279, April 2006. 5
- Mikhail Golosov, Narayana Kocherlakota, and Aleh Tsyvinski. Optimal indirect and capital taxation. *Review of Economic Studies*, 70(3):569–587, 2003. 2, 13, 19
- Mikhail Golosov, Aleh Tsyvinski, and Iván Werning. New dynamic public finance: A user’s guide. *NBER Macroeconomics Annual 2006*, 2006. 5
- Fatih Guvenen. Learning your earning: Are labor income shocks really very persistent? *American Economic Review*, 97(3):687–712, June 2007. 36
- Fatih Guvenen. An empirical investigation of labor income processes. *Review of Economic Dynamics*, 12(1):58–79, January 2009. 36
- John C. Hause. The fine structure of earnings and the on-the-job training hypothesis. *Econometrica*, 48:1013–1029, 1980. 36
- Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. Two views of inequality over the life-cycle. *Journal of the European Economic Association (Papers and Proceedings)*, 3:543–52, 2004. 33
- John Heaton and Deborah J Lucas. Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy*, 104(3):443–87, June 1996. 36
- Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5–6):953–969, 1993. 23, 31
- Kenneth L. Judd. Redistributive taxation in a perfect foresight model. *Journal of Public Economics*, 28:59–83, 1985. 2
- Narayana Kocherlakota. Zero expected wealth taxes: A mirrlees approach to dynamic optimal taxation. *Econometrica*, 73:1587–1621, 2005. 2
- Dirk Krueger and Fabrizio Perri. On the welfare consequences of the increase in inequality in the united states. In Mark Gertler and Kenneth Rogoff, editors, *NBER Macroeconomics Annual 2003*, pages 83–121. MIT Press, Cambridge, MA, 2004. 32
- Finn E. Kydland. A clarification: Using the growth model to account for fluctuations: Reply. *Carnegie-Rochester Conference Series on Public Policy*, 21:225–252, 1984. 36
- Ethan Ligon. Risk sharing and information in village economics. *Review of Economic Studies*, 65(4):847–64, 1998. 2

- Lee A. Lillard and Yoram A. Weiss. Components of variation in panel earnings data: American scientists, 1960-70. *Econometrica*, 47:437–454, 1979. 36
- David G. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, Inc, 1969. 22
- Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, March 2002. 44
- William P. Rogerson. Repeated moral hazard. *Econometrica*, 53(1):69–76, 1985. 2
- Robert Shimer and Ivan Werning. Liquidity and insurance for the unemployed. *American Economic Review*, 98(5):1922–42, December 2008. 5
- Daniel T. Slesnick and Aydogan Ulker. Inequality and the life-cycle: Consumption. Technical report, 2004. 33
- Stephen E. Spear and Sanjay Srivastava. On repeated moral hazard with discounting. *Review of Economic Studies*, 54(4):599–617, 1987. 23
- Nancy L. Stokey and Sergio Rebelo. Growth effects of flat-rate taxes. *Journal of Political Economy*, 103(3):519–50, 1995. 31
- Kjetil Storesletten, Chris I. Telmer, and Amir Yaron. Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717, 2004a. 28, 32, 36
- Kjetil Storesletten, Christopher I. Telmer, and Amir Yaron. Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633, April 2004b. 36
- Iván Werning. Optimal dynamic taxation. PhD Dissertation, University of Chicago, 2002. 2

